

The Price of Delay: Supply Chain Disruptions and Pricing Dynamics

Salomé Baslandze and Simon Fuchs

Working Paper 2025-8

September 2025

Abstract: We study the role of supply chain disruptions in shaping consumer prices, focusing on both firms' own import shocks and strategic responses to competitors' disruptions. Using a newly constructed microlevel dataset that links transaction-level US import data from bills of lading with high-frequency consumer prices and sales from a consumer panel, we develop a novel approach to estimate the price effects of cost shocks and product availability. Motivated by a model of delivery delays, cost shocks, and firm pricing, we implement a shift-share identification strategy based on delivery shortfalls, port congestion, and freight and import costs. We find sizable pass-through elasticities: firms raise prices in response to higher import costs and delivery delays, especially when disruptions persist. We also identify strategic pricing: firms—including non-importers—increase prices in response to competitors' supply chain disruptions. Using our estimates and back-of-the-envelope calculations from the model, we show that strategic interactions significantly amplified the direct effects of supply chain shocks on consumer prices during the pandemic.

JEL classification: E31, F14

Key words: supply chains, inflation, delivery delays, strategic interactions, pass-through, inventory

<https://doi.org/10.29338/wp2025-08>

The authors are grateful to KC Pringle and Michael Sparks for outstanding research assistance. They thank seminar and conference participants at the University of Mississippi, Monash University, ISOT, the SEM Conference (Georgia Tech), the Firm Dynamics Workshop (Federal Reserve Bank of Atlanta), the Virtual Seminar Series in Innovation and Growth, the Federal Reserve Bank of Minneapolis, SED, the BSE Summer Forum, Workshop on International Economic Networks (WIEN), Armenian Economic Association, and the System Committee on Macroeconomics for helpful feedback. They are especially indebted to their discussants, Davin Chor, Sebastian Heise, and Viacheslav Sheremirov, and to Leo Feler for support with the Numerator data. They also benefited from insightful comments by Pol Antràs, Michael McMahon, and Jon Willis. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Atlanta, the Federal Reserve System, or its staff.

Please address questions regarding content to Salomé Baslandze, Federal Reserve Bank of Atlanta and CEPR, baslandze.salome@gmail.com; or Simon Fuchs, Federal Reserve Bank of Atlanta, sfuchs.de@gmail.com.

Federal Reserve Bank of Atlanta working papers, including revised versions, are available on the Atlanta Fed's website at www.frbatlanta.org. Click "Publications" and then "Working Papers." To receive e-mail notifications about new papers, use frbatlanta.org/forms/subscribe.

1 Introduction

Throughout the pandemic, supply chain disruptions and their impact on inflation have been at the forefront of policy discussions. International trade and supply chains were disrupted in many different ways, ranging from backlogs at ports, heightened maritime shipping costs, production disruptions at the origin due to COVID restrictions, as well as increased demand shocks that put strains on limited production capacity. These various measures of disruptions are summarized by the NY Fed Global Supply Chain Pressure Index (GSCPI) depicted in blue in Figure 1. At the same time, US inflation has experienced historically high levels, as illustrated by CPI inflation depicted in Figure 1 in red. While the growing literature studies the implications of these supply chain disruptions,² it remains unclear what the firm-level pass-through of these disruptions to consumer prices is and what role strategic interactions play for the effect of shocks on consumer price dynamics.

To study these questions, we proceed in the following steps. First, we develop a motivating model of delivery delays, product availability, and firm pricing to illustrate—in a transparent way—how supply chain pressures translate into higher prices and to provide the estimating framework. Second, we assemble a large-scale micro-level dataset combining product-level prices with firm-level measures of delivery shortfalls, port congestions, and marginal cost shocks from unit import costs and freight costs. Using these data, we then estimate how firms’ own supply chain disruptions—delivery shortfalls, unit import cost shocks, and freight cost shocks—affect pricing using OLS and a shift-share instrumental variable approach. Next, we extend the framework to incorporate strategic interactions and estimate the effect of competitors’ disruptions on firm pricing, over and above the firm’s own shocks. Finally, we use the estimated objects in a simple accounting exercise to show how supply chain disruptions could shape aggregate price dynamics through direct cost effects, own availability effects, rivals’ availability effects, and the amplification generated by strategic feedback.


We begin by building a simple model that delivers an optimal inventory–pricing benchmark that links replenishment lead times and base-stock targets to product availability and markups. Product availability—the probability that an arriving customer is served—falls with longer lead times or lower base-stock and rises with faster replenishment. We show that lower availability reduces the effective price elasticity faced by the firm, which raises the optimal markup, while higher availability increases the effective elasticity and compresses markups. Consequently, supply chain shocks through longer delays and lower availability naturally result in higher prices.

We use our simple model to deliver reduced-form expressions for optimal price dynamics that split

²See, among others, [Ascari, Bonam and Smadu \(2024\)](#); [Bai et al. \(2024\)](#); [Finck, Klein and Tillmann \(2024\)](#); [Brancaccio, Kalouptsi and Papageorgiou \(2025\)](#).

observed price changes into (i) movements in marginal cost and (ii) movements in product availability, both filtered through the same pass-through parameter implied by the pricing rule. We map import delivery shortfalls to changes in availability and decompose marginal cost shocks into unit import and freight cost components. The simple model also disciplines the signs: higher availability raises the effective price elasticity and compresses markups, while increases in cost push prices up with a pass-through strictly between zero and one.

Figure 1. Supply Chain Pressure Index and Inflation



Notes: (A) Global Supply Chain Pressure Index (GSCPI, NY Fed). (B) CPI for All Urban Consumers: all items less food and energy, year-to-year change, from BLS. (FRED, CPILFESL)

To trace how upstream supply chain disruptions affect downstream consumer prices, we build a matched firm–product panel that combines two unusually granular sources: shipment-level import records, which reveal a firm’s supply pipeline at high frequency, and receipt-level consumer transactions, which record the prices actually paid by consumers for specific products over time.

On the upstream side, we use S&P Global Panjiva Bills of Lading (2007–2023), a transaction-level dataset of U.S. maritime imports reporting shipper–consignee pairs, HS descriptions, quantities/weights, and vessel arrival dates.³ Event-level time stamps and counterparty identifiers allow us to construct firm×product-code (HS) measures of delivery shortfalls, defined as deviations in k -month cumulative import volumes relative to the 2019 pre-pandemic benchmark, which we then aggregate to the firm level. We complement our measures of delivery shortfalls with firm exposure to port congestion (dwell times) by combining data on a firm’s port usage with U.S. port dwell time data from [Fuchs and Wong](#)

³Panjiva is based on the near-universe of waterborne bills of lading filed with U.S. Customs and Border Protection, providing named shipper–consignee identifiers and event-time stamps for individual shipments.

(2022).⁴ Finally, we assemble two model-motivated marginal-cost shifters: unit import costs and freight costs. Changes in import costs (unit values at the firm $\times$ HS-code level) and freight costs (at the port level) are aggregated using firm exposures to HS codes and ports to derive firm-specific shocks.

To measure downstream consumer prices, we use Numerator’s Consumer Panel (2019–2023), which records itemized receipts from both brick-and-mortar retailers and online. The data span a wide range of goods in the consumer basket, covering many product categories with detailed information on quantities, sales, producer names, brands, and purchase timestamps. We define products at the firm–brand–category level and track 12-month price changes as our baseline outcome. The receipt-level detail ensures we capture realized transaction prices, including promotions and coupons, across households, channels, and geographies. Aggregate spending in the panel closely follows Census retail sales, and the resulting inflation series comoves tightly with the CPI, underscoring the reliability of the data for analyzing consumer price dynamics.

We link consumer purchases from Numerator to international shipments in Panjiva data, creating a novel product-firm-level dataset that connects U.S. consumer prices to import activity. The matched sample covers more than 40% of all product-month observations in Numerator and about half of the total sales. Two descriptive patterns stand out: delivery shortfalls during 2020–2022 were large but highly uneven—both across categories and across firms within a category—and price changes display similarly wide dispersion, with the largest firms raising prices more on average and more diversified importers showing attenuated price growth.

Using these data, we estimate how supply-chain disruptions—delivery shortfalls, import unit cost shocks, and freight cost shocks—affect firm pricing. We regress 12-month product-level price changes on measures of firm-level shortfalls and marginal-cost shocks, controlling for firm and product category–time fixed effects that filter out the aggregate demand factors. Still, firm-level delivery shortfalls may be correlated with unobserved idiosyncratic demand shocks, biasing OLS estimates downward. To isolate the causal effect of availability disruptions, we implement a shift–share IV strategy that instruments delivery shortfalls with measures of delivery shortfall exposure and dwell-time change exposure based on firms’ pre-pandemic import portfolios. This design provides plausibly exogenous variation in supply conditions for firms within narrowly defined product markets facing the same aggregate demand shocks and yields, delivering a direct empirical counterpart to the theoretical pass-through equation.

We find that both own delivery shortfalls and marginal-cost shocks translate into sizable, but incomplete, changes in consumer prices. OLS estimates understate the effect of shortfalls, but instru-

⁴Similar data have also recently been used by [Brancaccio, Kalouptsi and Papageorgiou \(2024, 2025\)](#); [Bai et al. \(2024\)](#) to study supply chain disruptions, but to the best of our knowledge this is the first study that explicitly constructs and exploits heterogeneous incidence at the firm-level.

menting with exposure measures produces stable elasticities above 0.21. Delivery shocks (at 3-month moving averages) pass through contemporaneously, whereas import and freight cost shocks transmit with a short lag. Pass-through intensifies when disruptions persist, indicating that firms adjust prices more when they perceive shocks as durable. Heterogeneity analyses reveal that while durable goods experienced stronger price growth during the pandemic, this largely reflected demand pressures rather than a higher elasticity of supply-side pass-through. Overall, the results highlight that both availability constraints and cost shocks contributed meaningfully to consumer price inflation, with the IV strategy ensuring the estimates reflect supply-driven rather than demand-driven effects.

We next examine strategic interactions to assess whether firms' price responses to supply chain shocks were amplified by their rivals' disruptions. Extending the simple model to allow for strategic interactions ([Amiti, Itskhoki and Konings, 2019](#)) highlights how each firm sets prices not only based on its own costs and availability but also on competitors' conditions. Empirically, we augment our pass-through specification to include rivals' market-share-weighted delivery shortfalls and cost shocks (unit and freight), defining rivals within finely delineated product categories.

The results show that, beyond the direct effect of own disruptions, supply shocks also transmit to prices through strategic channels. Firms adjust prices upward in response to their competitors' disruptions, with the elasticity of rivals' shortfalls to own prices about half the size of the own-pass-through elasticity. Competitors' cost shocks likewise spill over, indicating that supply bottlenecks and cost pressures propagate through competitive interactions. Importantly, these effects extend to non-importers, who raise prices when importing rivals are hit, consistent with demand shifting toward unaffected sellers and reduced competitive discipline. Finally, both own and competitor pass-through are state-dependent: price responses are significantly larger in high-inflation sectors and during periods of elevated aggregate inflation. Together, these findings imply that supply chain disruptions can generate aggregate price pressures well beyond the directly affected firms, with amplification that is strongest precisely when inflation is already high.

Finally, we conduct a model-consistent accounting exercise (in an impulse-response spirit) that feeds aggregate shock series—delivery shortfalls, port congestion, and imported-input costs—into the estimated micro structure to trace their impact on prices. The decomposition shows that availability shocks, particularly those amplified by port dwell times and strategic complementarities, dominate the aggregate price surge dynamics. Strategic interactions magnify both own and rivals' disruptions, transmitting cost and availability shocks even to non-importers, with the strength of amplification shaped by market composition.

Literature. Our paper contributes to several areas of the literature. First, we speak to the pass-through literature for internationally sourcing firms by placing availability via delivery shortfalls—a reduced-form proxy for delivery delays and inventory constraints—on the same pricing margin as marginal costs. Cost-focused work shows that tariff cuts reduce marginal costs yet raise markups, implying incomplete pass-through (De Loecker et al., 2016); related evidence on energy shocks documents full pass-through for cost increases but only partial pass-through for cost reductions, with limited aggregate inflation effects given energy’s small cost share and substantial within-industry heterogeneity (Lafrogne-Joussier, Martin and Mejean, 2023). On the strategic side, Amiti, Itskhoki and Konings (2019) document strong strategic complementarities in price setting, and Albagli et al. (2025) show these complementarities can dominate cost forces and are state-dependent. We fold these insights into a unified specification that treats own costs, competitors’ prices, and—critically—own and rivals’ availability in parallel, thus isolating a distinct mechanism whereby supply-chain disruptions via delivery delays shift prices even conditional on costs. Our contribution is novel evidence on the significance of supply chain disruptions both for the direct - own - pass-through as well as strategic interactions between firms.

We also contribute to the literature that links inventories, delivery frictions, and pricing. Dynamic inventory models with monopolistic competition show that scarce stocks raise markups and slow pass-through, with delivery lags and fixed order costs generating state-dependent pricing and precautionary inventory behavior (Alessandria, Kaboski and Midrigan, 2010; Alessandria et al., 2023). Work on sourcing under stochastic delivery times emphasizes that binding inventory constraints tilt pricing toward precautionary markups and induce higher buffers when supply is slower or riskier (Carreras-Valle and Ferrari, 2025). Macro and retail studies with inventories and nominal frictions likewise produce real rigidities, inventory-driven sales, and state-dependent markups (Khan and Thomas, 2007; Kryvtsov and Midrigan, 2013; Aguirregabiria, 1999). Relative to these contributions, our approach is deliberately minimal and analytically transparent: it delivers closed-form objects that pin down an availability-adjusted markup and a sharp, estimable decomposition of price changes into cost and availability channels. This tractability directly motivates reduced-form empirical specifications for both own pass-through and strategic exposure to rivals’ conditions, and it provides a clean bridge to frameworks that measure strategic complementarities in price setting—most notably extending Amiti, Itskhoki and Konings (2019) by bringing firm and rival availability into the same empirical mapping as costs and competitor prices.

Finally, we contribute to a string of recent research that links pandemic-era bottlenecks to inflation and macro dynamics. Survey evidence shows disruptions shifting expected unit costs (Meyer, Prescott and Sheng, 2023); New Keynesian analyses with binding capacity constraints trace upward shifts in

Phillips curves and amplified goods-sector inflation (Comin, Johnson and Jones, 2023); indices based on container-ship movements and DSGE exercises attribute a sizable 2021 inflation impulse to supply-chain shocks and study monetary-policy interactions (Bai et al., 2024; Amiti et al., 2024); and high-frequency micro evidence documents how pandemic-induced stockouts transmitted into temporary but significant inflationary pressures (Cavallo and Kryvtsov, 2023). Related work emphasizes the central role of transportation rigidities—ships and ports—in propagating shocks to trade costs and inflation (Brancaccio, Kalouptsi and Papageorgiou, 2025). At the micro level, Liu, Smirnyagin and Tsyvinski (2024) show that delivery lapses depress supplier balance sheets and firm performance without tracing consumer-price pass-through, while Borusyak and Jaravel (2021) study tariff pass-through but not pandemic-related delays. Our contribution is to provide causal evidence—using unique micro-level data that link firm-specific delivery shortfalls to item-level prices—on how supply-chain disruptions move prices and how strategic complementarities across firms amplify these effects. We isolate delay-specific shocks from input-cost movements and pair the evidence with a tractable model-to-data mapping that delivers closed-form pass-through objects and a transparent decomposition of price changes into cost and availability channels, quantifying both own pass-through and strategic exposure to rivals’ conditions.

The remainder of the paper is structured as follows. Section 2 introduces a motivating model. Section 3 describes the data construction and provides summary statistics. Section 4 analyzes the pass-through of both own and competitors’ supply chain disruptions to consumer prices. Section 5 presents a simple accounting exercise to quantify potential aggregate implications of supply chain disruption shocks. Section 6 concludes.

2 Theory

In this section, we lay out a simple, transparent model that connects supply-chain conditions to firms’ pricing and inventory choices. The model provides economic intuition and motivates our empirical design. A downstream firm facing flow demand and stochastic replenishment chooses a posted price and a base-stock target. Under a base-stock policy, we obtain closed-form expressions for the stock-in probability and a generalized Lerner rule in which an availability wedge (driven by congestion in replenishment) raises markups when stock is scarce. Linearizing the pricing condition delivers a two-term decomposition of price changes into a marginal-cost component and an availability component—a structure we take directly to the data. The derivations and notation are collected in Appendix A.1; in Section 4.2 we extend the model to allow for strategic interactions between firms.

2.1 A simple model of pricing, availability, and lead times

We study a single downstream firm that *procures* a ready-made product from an upstream supplier rather than producing it in-house.⁵ The firm chooses a posted (log) price p and a base-stock target τ , our *availability shifter* that summarizes order size, review frequency, expediting, and sourcing intensity. Operationally, a base-stock (order-up-to) rule aims to restore on-hand inventory to τ whenever a delivery arrives; higher τ raises the chance that an arriving customer is served but increases holding cost, so τ is the reduced-form buffer-capacity lever given the upstream replenishment process.⁶

Customer demand arrives as a flow at rate $\lambda(p)$, decreasing in price. Replenishment of the firm's inventory occurs through the supplier with *stochastic* delivery times summarized by a *replenishment speed* $\mu > 0$: higher μ means faster deliveries (mean lead time $1/\mu$). Only demand arrivals that find stock on hand are served. Let p and c denote *log* price and *log* marginal cost, with corresponding *levels* $P := e^p$ and $MC := e^c$, and let $h > 0$ be the holding-cost rate per unit of average on-hand inventory $\mathbb{E}[I(\tau, \lambda(p), \mu)]$. Finally, we are making the assumption that customers who arrive when no stock is hand are not being served, an assumption commonly called the *lost-sales setting*. Focusing on steady-state flow payoffs in a lost-sales setting, the firm's static profit rate is⁷

$$\Pi(p, \tau) = (P - MC) \lambda(p) s(\tau, \lambda(p), \mu) - h \mathbb{E}[I(\tau, \lambda(p), \mu)],$$

where $s(\tau, \lambda, \mu) \in [0, 1]$ is the acceptance (stock-in) fraction—the probability that a demand arrival is served.

To obtain closed-form objects that map cleanly to the data, we adopt two standard assumptions. (i) Flow demand is log-linear, $\lambda(p) = \Lambda e^{-\sigma p}$, so $\sigma > 0$ is the demand semi-elasticity, and $\Lambda > 0$ collects demand shifters. (ii) Replenishment lead times are exponential with rate $\mu > 0$ (mean $1/\mu$), and the

⁵Although we maintain this interpretation, the model can readily be adapted to view firms as receiving intermediate goods deliveries—an interpretation that maps naturally to our import data discussed later.

⁶We deliberately abstract from several operational features to keep the mapping transparent: a single product under continuous review with lost sales (rather than backlogging), no fixed order or setup costs or minimum order quantities (which would induce (s, S) policies), and no storage limits, perishability, or order crossover. We also treat the replenishment speed μ as within-period exogenous. As we will show below, these simplifications deliver the closed-form acceptance probability $s(\tau, \lambda, \mu)$ and a clean link between pricing and inventory that we can take to the data. Formal details and variants are provided in Appendix A.1.

⁷Lemma 1 in Appendix A.1 shows that the profit rate is an equivalent objective to long-run average profit under the memoryless property of the exponential distribution, (A23)–(A25).

firm follows a base-stock rule at level $\tau \in \mathbb{N}$.⁸ Define

$$r \equiv \frac{\lambda}{\lambda + \mu} \in (0, 1), \quad s(\tau, \lambda, \mu) = 1 - r^\tau,$$

so r measures how fast demand arrives relative to replenishment (“congestion”), and one can show that under the assumptions above s is the probability an arriving customer finds stock on hand - our closed-form stock-in probability.⁹ It is also convenient to track how demand pressure erodes availability via

$$\kappa(\tau, \lambda, \mu) \equiv \lambda \frac{\partial s}{\partial \lambda} = -\tau(1 - r)r^\tau < 0,$$

an *availability adjustment* that will enter the pricing condition. Intuitively, r summarizes congestion, s summarizes how often customers are served, and κ summarizes how sensitive that service probability is to demand pressure—three workhorse objects that link lead times and inventory choices to prices and that carry through to the empirical mapping below.

Under these stochastic assumption, we have derived a mildly adjusted standard firm optimization problem. The firm chooses a profit maximizing price, but in our setting also chooses a base-stock to maintain product availability and internalizes the effect pricing has on availability. Given this setup, we are now in a position to characterize the firm’s optimal joint pricing and inventory choice. The following proposition summarizes the firm’s optimal choice under the additional assumption that pricing is constant within a replenishment cycle.¹⁰

Proposition 1 (Optimal pricing and inventory in the simple model). *For any τ , the profit-maximizing*

⁸A realistic microfoundation for exponential lead times comes from the assumptions that upstream suppliers fill orders subject to capacity constraints, with orders that are waiting to be served forming a queue. Formally, we can be generic about both the arrival and servicing process of that upstream queue as long as we accept an asymptotic characterization of the queue under heavy traffic. In queuing notation, we consider a upstream G/G/1 (single-server queue with interarrival and servicing times having an arbitrary distribution) queue in heavy traffic: the stationary lead time is approximately $\text{Exp}(\mu_{\text{HT}})$ with μ_{HT} proportional to system slack and dampened by variability; see Appendix A.1, Lemma 2.

⁹Appendix A.1 first derives the stock-in probability from the exponential arrival processes for customers and upstream suppliers and then furthermore shows that the arrival probability for any individual customers, s , equals the time-average availability use in the objective function, the profit rate (Lemma 1).

¹⁰We assume prices are fixed within a replenishment cycle. This is consistent with micro evidence on infrequent price adjustment and with standard pricing frictions. For U.S. consumer prices, Nakamura and Steinsson (2008) show that once sales are netted out, the median frequency of nonsale price changes is on the order of 9–12% per month (implying median durations of roughly 8–11 months), whereas including sales yields much shorter durations. For the euro area, Alvarez et al. (2006) document average monthly frequencies near 15% (implying typical durations close to a year), again indicating that many prices remain unchanged for several months at a time. These facts make it empirically plausible that a firm does not reprice within a single replenishment cycle. The assumption also buys tractability: it delivers closed-form characterizations for availability, markups, and pass-through that map cleanly to reduced-form specifications.

price satisfies the generalized Lerner condition in levels:

$$\frac{P^* - MC}{P^*} = \frac{1}{\sigma \left(1 + \frac{\kappa}{s}\right)} \quad \left(\text{equivalently } p^* = c - \ln \left[1 - \frac{1}{\sigma(1 + \kappa/s)}\right]\right),$$

with $s = 1 - r^\tau$, $\kappa = -\tau(1 - r)r^\tau$, $\lambda = \Lambda e^{-\sigma p^*}$, and $r = \lambda/(\lambda + \mu)$. If τ is chosen discretely, an interior base-stock τ^* solves the exact marginal condition

$$(P^* - MC) \lambda r^{\tau^*} (1 - r) = h(1 - r^{\tau^*+1}),$$


with $\lambda = \Lambda e^{-\sigma p^*}$. Under the continuous relaxation of τ , this first-order condition admits a unique solution with $r^{\tau^*} \in (0, 1)$.

These results have a simple interpretation. The pricing rule is a generalized Lerner condition in which the usual elasticity σ is adjusted by the *availability term* $1 + \kappa/s$. Because $\kappa/s < 0$ whenever stockouts can occur, the effective elasticity $\sigma(1 + \kappa/s)$ is below σ , implying a markup fraction $(P^* - MC)/P^*$ above the frictionless benchmark $1/\sigma$. The inventory condition trades off the marginal contribution from an additional unit of buffer against its holding cost: an extra unit pays off exactly when, before the next delivery, the next τ superposed events are $\tau - 1$ demand arrivals followed by a delivery, an event with probability $r^\tau(1 - r)$, yielding marginal revenue $(P^* - MC)\lambda r^\tau(1 - r)$ that is set equal to $h(1 - r^{\tau+1})$. As replenishment becomes faster, that event becomes less likely, the shadow value of inventory falls, and the optimal base-stock τ^* declines. Taken together, higher μ improves availability through two channels: it directly raises the acceptance probability s and, by pushing κ/s toward zero, increases the effective elasticity and compresses the markup; both effects reduce the incentive to hold inventory. Panels 2a–2b provide a compact visualization of the two optimality margins in Proposition 1. Panel 2a plots the inventory condition: the optimal base-stock is where the marginal revenue from an extra buffer unit intersects the marginal holding-cost schedule; when delivery slows (lower μ), the marginal-revenue curve shifts up and the optimal τ rises. Panel 2b depicts the pricing condition in accepted units: the availability wedge shifts marginal revenue below the frictionless benchmark, so tighter availability (lower s) raises the markup and reduces the accepted flow; as availability improves, the outcome moves toward the frictionless point.

Corollary 1 (Comparative statics with respect to lead times μ). *Holding primitives (Λ, σ, c, h) fixed,*

$$\partial_{\ln \mu} \ln s = \frac{\tau r^\tau (1 - r)}{1 - r^\tau} > 0, \quad \partial_{\ln \mu} P^*(\mu) < 0, \quad \partial_{\ln \mu} \tau^*(\mu) < 0.$$

Figure 2. Inventory management with delivery delays



Notes. *Panel (a) — Inventory choice.* On the horizontal axis: base-stock (availability shifter) τ . On the vertical axis: marginal value/cost per unit time. $MC^{inv}(\tau)$ is the marginal holding-cost schedule. $MR^{inv}(\tau, \mu)$ is the marginal revenue from one more unit of buffer. When replenishment slows ($\mu^2 < \mu^1$), $MR^{inv}(\tau, \mu)$ shifts up—each unit of buffer is more valuable—so the optimal base-stock rises: $\tau^*(\mu^2) > \tau^*(\mu^1)$ where $MR^{inv} = MC^{inv}$. The shaded area illustrates inventory surplus at the higher μ ; it expands when delivery slows. *Panel (b) — Price choice in accepted units.* On the vertical axis: price p . On the horizontal axis: accepted flow q^A . Inverse demand is $p = D(q^A)$. The availability-adjusted marginal-revenue curve $MR^s(q^A)$ lies below the frictionless benchmark $MR^0(q^A)$ (the case $s \equiv 1$). Price is set where MR^s meets marginal cost $MC(q^A)$. Tighter availability (lower s) shifts MR^s down, yielding higher prices and lower accepted quantities ($q^{A,\tau}, p^\tau$); as s improves, the outcome moves toward the frictionless point (q^{A,p^c}, p^{p^c}). The ordering $p^{p^c} < p^m < p^\tau$ visualizes the generalized-Lerner logic: weaker availability lowers effective elasticity and raises the markup.

In the fast-replenishment limit $\mu \rightarrow \infty$, we have $r \rightarrow 0$ and $s \rightarrow 1$; by the Lerner rule the markup fraction approaches $1/\sigma$, so the optimal level price tends to $MC \cdot \sigma/(\sigma - 1)$ and the optimal log price tends to $c + \ln(\sigma/(\sigma - 1))$. Conversely, in the slow-replenishment limit $\mu \rightarrow 0$, we have $1 - r \approx \mu/(\lambda + \mu)$ and hence $s \approx \tau \mu/(\lambda + \mu)$; in this region the price response becomes locally flat in μ with $\partial_{\ln \mu} p^* \rightarrow 0$.

Corollary 1 says, in plain terms, that faster replenishment makes more arrivals get served, scarcity fades, and the firm faces a “softer” constraint—so markups and prices fall. At the same time, an extra unit of buffer is less valuable (stockouts are less likely before the next delivery), so the optimal base-stock declines. These forces bite most in the middle of the state space—when the system is neither almost always stocked nor almost always empty—because small improvements in delivery speed then move both availability and prices the most. At the extremes, effects are muted: with near-instant delivery the frictionless markup prevails; with very slow delivery availability is so low that further slowdowns barely move pricing. To see an illustration of the comparative statics see Figure 2: in panel (a), faster delivery flattens the marginal-revenue curve and lowers the optimal buffer.

To recap: Our simple model delivers two workhorse objects in closed form—a stock-in (availability) probability and an availability-adjusted markup rule—and uses them to generate intuitive comparative statics: faster replenishment raises availability, compresses markups, and lowers the value of holding inventory, while slower or riskier replenishment does the opposite. These closed-form objects also yield a ready-to-estimate linear decomposition of price changes into cost and availability components, which

underpins our empirical design. Relative to existing work on joint pricing and inventory, which typically solves dynamic models numerically, our approach keeps the same core mechanism but in a tractable, closed-form environment. For example, [Alessandria, Kaboski and Midrigan \(2010\)](#) study a dynamic importer with (S, s) ordering and delivery lags that link markups to the shadow value of inventories, while [Alessandria et al. \(2023\)](#) show in a heterogeneous-firm GE setting that uncertain shipping delays are contractionary and raise prices via stockouts; our scarcity wedge captures the same channel in reduced form. Likewise, [Carreras-Valle and Ferrari \(2025\)](#) obtain precautionary markups under stochastic delivery times, mirroring our availability channel.¹¹ Our contribution is to isolate these forces in a mildly extended standard firm optimization problem, delivering analytic expressions that map directly to reduced-form pass-through—both own and strategic—without solving a dynamic program. We next turn towards what the model implies for our empirical design.

2.2 From Theory to Empirics

By Proposition 1, the optimal price is pinned down by two ingredients: marginal cost and an availability-adjusted scarcity term (the κ/s wedge) that compresses the effective demand elasticity when stockouts are likely. Linearizing around an operating point yields¹²

$$dp = \underbrace{\frac{1}{1+\Gamma}}_{:=\alpha \in (0,1)} dmc + \underbrace{\frac{\Lambda}{1+\Gamma}}_{:=\beta_s < 0} d \ln s + \varepsilon, \quad (\text{E2})$$

where α is *own pass-through* and β_s the *availability channel*. The intuition is simple: any price change feeds back into scarcity—raising p reduces arrivals, improves availability, and softens the desired markup—so *both* cost and availability shocks are damped by the same factor $1/(1 + \Gamma)$. When this scarcity feedback is weak ($\Gamma \approx 0$), pass-through is close to one; when it is strong, pass-through is muted. Because $\beta_s = \alpha \Lambda$ with $\Lambda < 0$, better availability raises effective elasticity and lowers the desired markup, and the resulting price change is scaled by the same α . In short, costs shift the cost side and

¹¹Furthermore, classic inventory macro models ([Khan and Thomas, 2007](#); [Kryvtsov and Midrigan, 2013](#)) imply similar comparative statics—faster replenishment compresses markups and weakens buffer motives—which we quantify in a tractable flow-demand, exponential-lead-time environment. See Appendix A.3 for an explicit comparison with [Kryvtsov and Midrigan \(2013\)](#).

¹²Where we define the markup mapping $\mathcal{M}(p, \ln s) := \sigma(1 + z(p, \ln s))[\sigma(1 + z(p, \ln s))]^{-1}$ so that $\Gamma := -\partial \mathcal{M} / \partial p > 0$ and $\Lambda := \partial \mathcal{M} / \partial \ln s < 0$. Start from the Lerner rule in Proposition 1 and write the price–cost gap as a markup mapping

$$\mathcal{M}(p, \ln s) := -\ln\left(1 - \frac{1}{\sigma[1+z(p, \ln s)]}\right), \quad z := \kappa/s.$$

Then the optimality condition is $F(p, c, \ln s) := p - c - \mathcal{M}(p, \ln s) = 0$. Totally differentiating around an operating point gives $(1 + \Gamma) dp = dc + \Lambda d \ln s$, where $\Gamma := -\partial_p \mathcal{M}(p, \ln s) \geq 0$ and $\Lambda := \partial_{\ln s} \mathcal{M}(p, \ln s) \leq 0$. Rearranging yields the linear approximation $dp = \alpha dc + \beta_s d \ln s$ with $\alpha = 1/(1 + \Gamma)$ and $\beta_s = \Lambda/(1 + \Gamma)$. Replacing dc by dmc (the log marginal–cost change) gives (E2). Detailed derivations in Appendix A.1.

availability shifts the markup side, but a single attenuation parameter α governs how quickly prices move.

From theory to measurement. To take the model to the data, we use two observables: (i) unit–import and freight price indices that move marginal cost, and (ii) a *delivery shortfall* index S_{it} that proxies for availability. Intuitively, shortfalls in imports slow replenishment, draw down inventories I , and reduce the stock–in probability s , which raises prices through the availability term in Proposition 1. these are my comments Locally, we summarize this link by

$$\Delta \ln s_{it} \approx -\phi_i S_{it}, \quad \phi_i := \underbrace{\frac{\partial \ln s}{\partial \ln I}}_{\eta_{sI} > 0} \times \underbrace{\frac{\partial \ln I}{\partial \ln M}}_{\eta_{IM} \text{ (import dependence)}},$$

so the impact of a given shortfall is larger when availability is more sensitive to inventories and when inventories depend more on inbound flows. On the cost side, we decompose marginal–cost changes as $\Delta mc_{it} \approx \theta_{Mi} \Delta \ln P_{M,t} + \theta_{Fi} \Delta \ln F_t$, where θ_{Mi} and θ_{Fi} are product–level output elasticities. Combining these ingredients with the linear pricing response $dp = \alpha dmc + \beta_s d \ln s$ yields the estimating equation

$$\Delta p_{it} = \alpha(\theta_{Mi} \Delta \ln P_{M,t} + \theta_{Fi} \Delta \ln F_t) + (-\beta_s \phi_i) S_{it} + \text{FE} + \varepsilon_{it}, \quad (\text{E3})$$

so the coefficients on the cost series recover $\alpha\theta_{Mi}$ and $\alpha\theta_{Fi}$, while the shortfall coefficient $(-\beta_s \phi_i)$ is larger exactly where import flows are a more important source of inventory. (The precise construction of S_{it} is deferred to the data section.)

3 Data and Measurement

We construct a large-scale, micro-level dataset that links consumer prices to firms’ import activity in order to measure how delivery shortfalls and cost shocks pass through to prices. This linkage is achieved by combining two rich datasets: a unique consumer panel from Numerator,¹³ which provides detailed product-level price and sales data, and shipment-level Bills of Lading (BoL) data from Panjiva, which cover nearly the universe of U.S. imports and contain detailed economic and logistical information on each shipment. In what follows, we describe each dataset in detail and outline our construction of measures for price changes, delivery shortfalls, and marginal cost shocks.

¹³We thank Leo Feler and the data team at Numerator for providing us with data access and invaluable support.

3.1 Numerator Data and Price Measurement

We use consumer panel data from Numerator for the period 2019-2023 to infer prices and sales of various products by firms.¹⁴ Numerator is a marketing research company that collects data on consumer purchases in brick-and-mortar stores and online. Users of the Numerator app participate by either taking pictures of their receipts or passively sharing their digital purchases by allowing the app to track their online activity. The closest counterpart to these data is the widely used Nielsen Homescan (HMS) data from the Kilts Center.¹⁵ For our purposes, Numerator’s advantages over HMS include its coverage of the recent pandemic period and its inclusion of e-commerce, whose share of total retail sales has substantially increased since the pandemic.¹⁶ Appendix Section B.1 provides more details about the data.

The Numerator dataset contains over one billion receipts from 3.9 million users, forming the basis of a *static panel* of 450,000 panelists (200,000 monthly) who report purchases for at least 12 consecutive months and are representative of the U.S. population along key Census demographics. For each static panelist, the data include demographic weights (*demo_weight*) for alignment with gender, age, ethnicity, household income, household size, presence of children, and census division, and projection factors (*national_factor* and *trend_factor*) to scale from sample households to the U.S. population. We use the combined weights and factors in all analyses of quantities and sales.

Each shopping receipt recorded in the Numerator data provides detailed information about the purchased basket, including item descriptions, quantities purchased, prices, and details about the stores (such as name and address) where the purchases occurred. Important for our analysis is the information on the purchased items. Where available, Numerator item IDs are linked with UPC/GTIN codes, brands, and producer names.¹⁷ Items are classified into different categories with varying levels of aggregation. We exclude purchases that lack classification or are categorized as restaurants, non-items, or intermediate categories, focusing on item IDs with non-missing brand or manufacturer information.

The comparisons of Numerator data with other official statistics suggest that the data provide comprehensive coverage of product sales and consumer behavior in the U.S. First, the demographics of static panelists are similar to those reported in the 2019 Census. Appendix Figure A.2 compares demographics of the static panelists with the US Census and shows how the sample is very close to the representative

¹⁴Analysis is based on the April 19, 2024, Numerator Data Delivery.

¹⁵The Numerator data have only recently begun to be used in academic research, having been described and benchmarked against similar datasets in prior studies (He and Su, 2023; Hacıoğlu Hoke, Feler and Chylak, 2024; Hristakeva, Liaukonyte and Feler, 2024) and more recently applied in Baslandze et al. (2025).

¹⁶Approximately 90% of the transactions are from offline purchases, while the remaining transactions come from digital sites such as Amazon.com.

¹⁷Many products or services are not associated with UPC codes, such as restaurant or entertainment purchases.

US household in terms of age, region, education, employment, ethnicity, gender, number of children, marital status, and income (even without applying demographic weights). Second, changes in total sales across different purchase categories (such as apparel, electronics, health and beauty, and food) closely track the sales changes observed in U.S. Census Monthly Retail Sales data (Numerator, 2023). Most importantly for our analysis, monthly price changes in Numerator closely align with monthly price changes reported by the BLS for various purchase categories, as shown below.

For our baseline analysis, we aggregate 54 million item IDs reported in Numerator and define a *product* as the unique combination of the firm, brand, and product category. For example, a product might be a Hasbro (firm) producing a Peppa Pig (brand) 3-D puzzle (product category). This approach yields 645,890 distinct products and 77,395 brands sold by 9,458 firms across 4,000 different product categories.¹⁸

For each item and month, we compute the average price as weighted total sales divided by weighted total quantity, using all transactions from static panelists. *Product-level price* $Price_{p,t}$ — with p being a product and t being a month— is then obtained as a quantity-weighted average of item-level prices.

We construct the *product-level price change* as a 12-month smoothed price change for product p in month t , defined as:

$$\Delta_{12} Price_{p,t} = \frac{Price_{p,t} - Price_{p,t-12}}{(Price_{p,t} + Price_{p,t-12})/2}.$$

This measure, inspired by [Davis, Haltiwanger and Schuh \(1998\)](#), reduces the influence of outliers and captures relative price changes over a one-year horizon, allowing for robust analysis of pricing dynamics.¹⁹ To further reduce outliers and account for data noise, we trim the top and bottom 1% of 3, 6, and 12-month price changes (including missings) and drop all products with fewer than five observations over time and having fewer than five receipts collected in a month.


Numerator data offer a reliable measure of overall consumer price inflation, capturing the broader inflationary pressures that emerged across the economy during the pandemic. This is evident from the close alignment between price changes in the Numerator and official CPI data, both of which peak in mid-2022. Figure 3 illustrates this pattern by comparing 12-month inflation measures from Numerator with 12-month Urban CPI and Urban Commodity CPI from official statistics.

Finally, the availability of micro-level product price data reveals substantial heterogeneity in price changes over time— both across categories and, within categories, across firms— especially during the pandemic (Appendix Figure A.3). The dispersion is large enough that regressions of price growth or

¹⁸Other examples of products include *Hewlett Packard/Hewlett Packard/Calculators*, *Revlon Consumer Products/Revlon/Flat Irons*, *American Italian Pasta Company/Mueller's/Lasagna*, and *Loreal/Cerave/Bar Soaps*.

¹⁹We verify the robustness of our results by comparing this smoothed measure with an alternative definition in log changes, confirming that our findings are consistent across different specifications.

Figure 3. Numerator Price Change vs CPI



Notes: The solid red line illustrates the 12-month aggregate price changes derived from Numerator data. We compute mean product-level price changes within each category and then aggregate these category-level changes using product category sales weights. The dashed green and blue lines represent the Urban CPI for all items (FRED series: CPIAUCSL) and the Urban CPI for commodities (FRED series: CUSR0000SAC), respectively.

shortfalls on product category and time fixed effects yield R^2 values below 10%, providing substantial variation for our estimation.

3.2 Panjiva Bill of Lading Data (BoL) and Delivery Shortfalls

The analysis relies on bill of lading (BoL) data to infer import quantities and delivery shortfalls. A BoL is a legal document between the shipper and carrier required for the shipment of goods, detailing the commodity, shipping and receiving entities, vessel identifier, port of entry, payment terms, and other logistics details. These data are derived from images of contracts processed from U.S. ports and customs, providing a unique, high-frequency view of trade relationships at the transaction and firm level. BoL data are particularly valuable for capturing detailed maritime trade and logistics information in a timely manner, making them well-suited for analyzing international delivery delays.²⁰

We use BoL data from S&P Panjiva, which provide comprehensive records of U.S. maritime imports since 2007, covering over a billion shipments from 17 countries (see [Flaen et al. \(2023\)](#) for a detailed description, comparisons with Census data, and discussion of limitations).²¹ Because BoL does not report reliable import values, we use public Census data from U.S. Trade Online to convert BoL import volumes (in TEU) into values. This is done by computing HS2-specific 2019 volume-to-value ratios and

²⁰In 2020, maritime transport accounted for 43% of U.S. imports by value and 60% by weight (Bureau of Transportation Statistics).

²¹Since these are maritime shipments, coverage from Mexico and Canada is very limited.

applying them to BoL shipment volumes at the firm–HS2–month level. Appendix Section B.2 provides more details about the BoL data, cleaning procedure, and construction.

Our final dataset covers U.S. imports from 2019–2023, aggregated to a firm–HS2–month panel that includes shipment volumes and values, baseline (2019) benchmarks, and moving-cumulative measures of import volumes. These data allow us to measure deviations from typical shipment volumes to detect delivery delays, which we describe below. In addition, we combine these data with port congestion and shipping delay indicators derived from Automatic Identification System (AIS) data. These complementary sources, which capture vessel trajectories and port dwell times, help distinguish between delays arising from port congestion and those specific to the shipment’s origin.

Delivery Shortfalls and Instruments Construction

Delivery Shortfalls. Inspired by our discussion in Section 2, the core measure of interest, delivery shortfall, is constructed to quantify the deviations in k -month cumulative import volumes relative to a baseline year, capturing the impact of delivery delays at the firm and product level. Specifically, monthly delivery shortfalls are calculated as a smooth percent deviation of a k -month cumulative imports in a specific month from their 2019 levels in the same month. Hence, the year 2019 serves as a pre-disruption benchmark. Deviations from that benchmark “schedule” will proxy for inventory disruptions of the firm. The delivery shortfall at the firm-HS-code level is defined mathematically as:

$$\text{Delivery Shortfall}_{f,h,t(k)} = -\frac{\sum_{m=t-k}^t D_{f,h,m} - \sum_{m=t-k,2019}^{t,2019} D_{f,h,m}}{\left(\sum_{m=t-k}^t D_{f,h,m} + \sum_{m=t-k,2019}^{t,2019} D_{f,h,m} \right) / 2},$$

where $D_{f,h,m}$ represents the observed imports for firm f and product code (HS) h in month m .

This firm-HS-code-level measure is then aggregated to obtain a firm-level delivery shortfall by weighting the HS-code shortfalls by their share in total imports in 2019 (a baseline year before disruption):

$$\text{Delivery Shortfall}_{f,t(k)} = \sum_h \omega_{fh}^{2019} \times \text{Delivery Shortfall}_{f,h,t(k)},$$

where ω_{fh}^{2019} represents the share of each HS-code h in the total imports of firm f in 2019. This aggregation captures the overall impact of delivery disruptions at the firm level, accounting for the relative importance of each product category in the firm’s baseline portfolio.

Our delivery shortfall measures are constructed using a stable sample of firms that had some import

activity in 2019 and maintained an import presence in each subsequent year. This approach reduces data noise and focuses the analysis on firms that consistently rely on imports.²² For this sample, we identify each firm’s portfolio of HS codes and their import schedules in 2019 and track firm-HS-code (k -month cumulative) deliveries over time (months) relative to the benchmark in 2019 (including zeros).²³ Finally, we normalize the shortfall series to January 2020.

We use $k = 3$ as our baseline shortfall measure, but Figure 4 also shows how delivery shortfalls evolve under alternative moving averages ($k = 2$ and $k = 6$). These shortfalls capture supply chain disruptions, with larger positive deviations indicating more severe delays in the flow of goods. A sharp increase in 2020, relative to the 2019 benchmark, reflects the unprecedented disruptions triggered by pandemic-related lockdowns and labor shortages. Shortfalls rebounded in 2021, though the series continues to exhibit fluctuations. As expected, longer moving averages smooth out short-term volatility and result in less variable trends. Appendix Figure A.6 shows a wide heterogeneity in import shortfalls across different HS codes. We observe a long right tail of products facing severe import shortfalls throughout the recent period. This evidence highlights that supply chain challenges did not impact all sectors equally, with some industries facing significantly more severe constraints than others.

Figure 4. Average Delivery Shortfall Over Time



Notes: Weighted firm-level shortfalls over time. Panjiva sample. Weights are firm-level total annual imports.

The approach exploits the granularity of BoL data to build a proxy for delivery delays by comparing

²²There is substantial entry and exit in the BoL data, as shown in Appendix Figure A.5, which plots total import volumes for the sample of firms with any deliveries in 2019—comparable to Figure A.4. While focusing on a stable sample of importers helps reduce data noise, it also drops those firms that experienced extreme negative events and exited the market, potentially leading to underestimating the negative effects of supply chain disruptions.

²³This implies that entry of new HS codes within a firm over time obtains zero weights in these calculations, so that we trace a deviation relative to the 2019 benchmark.

current import volumes with pre-disruption levels. The shortfall measure captures deviations driven by shipping disruptions, production delays, and other logistical frictions, providing a direct reflection of firm-level supply chain shocks. However, we note that the change in delivery schedule can also be caused by aggregate or individual demand factors. In their raw form, delivery shortfall measures do not filter out these demand factors. In the regression analysis, we will include market-related trends to filter out aggregate demand factors and use the IV approach to filter out the individual demand forces.

Instrument 1: Delivery Shortfall Exposure. We construct two instruments for delivery shortfalls: Delivery Shortfall Exposure $_{f,t(k)}$ and Δ Dwell Exposure $_{f,t(k)}$. We define a firm’s exposure to delivery shortfalls as follows:

$$\text{Delivery Shortfall Exposure}_{f,t(k)} \equiv \sum_{h \in S^{HS2}} \omega_{fh}^{2019} \times \text{Delivery Shortfall}_{h,t(k)-f},$$

where $h \in S^{HS2}$ represents HS2 product codes, Delivery Shortfall $_{h,t(k)-f}$ denotes the HS code-level aggregate shortfall in a particular month(-year) t relative to that month in 2019, excluding the focal firm f itself (leavout shifts). The weight ω_{fh}^{2019} is the pre-determined share of firm f ’s imports in HS code h as of 2019, before major supply chain disruptions occurred. This measure is a standard shift–share instrument that exploits variation in firms’ import shares across HS codes, even within the same product categories—an identifying assumption we validate in the data.


Instrument 2: Port Congestion Exposure. Next, we construct a shift-share instrument for port congestion exposure using dwell-time variation across ports. The idea is that the longer the dwell times due to port congestion, the more likely a delivery shortfall is to occur. This instrument leverages exogenous variation in port-level congestion, derived from detailed port dwell time data, combined with firm-specific historical port usage patterns. The port dwell time data itself has been constructed from high-frequency Automatic Identification System (AIS) vessel traffic information which provides high-frequency location information of individual vessels at U.S. ports (Fuchs and Wong, 2022). Dwell time is then defined as the duration a vessel spends moored at a pier with zero speed. Crucially, to isolate port-specific conditions from variations due to the mix of ships calling, these raw dwell times (in logs) are residualized by controlling for ship characteristics, particularly size (e.g., gross tonnage). This yields a measure of port-level delay (in logs)(Residualized Dwell Time $_{pt}$) for port p at time (year-month) t that is net of ship-specific factors.

We combine this port-level “shift” measure with firm-specific “shares” derived from BoL data. The BoL data provides granular information on import shipments at the firm level, including the port of

entry. From this, we calculate each firm f 's historical reliance on different ports, measured as the share of its total imports that arrived through a specific port p during a pre-determined base period (ω_{fp}^{2019}). Using 2019 as a base period ensures these shares reflect established logistical relationships. The firm-level shift-share instrument, representing firm f 's exposure to port congestion, is then constructed by weighting the port-specific residualized dwell time shifts by the firm's historical port shares, summing across all ports:

$$\Delta \text{Dwell Exposure}_{ft} \equiv \sum_p \omega_{fp}^{2019} \times \Delta \text{Residualized Dwell Time}_{pt}.$$

Figure 5. Average Change in Dwell Time Exposure Over Time




Notes: Weighted firm-level change in dwell time exposure over time. Panjiva sample. Weights are firm-level total annual imports.

This instrument is designed to capture the variation in a firm's exposure to import delays driven by congestion at the specific ports it historically uses, isolating it from firm-specific demand or supply shocks. The validity of the instrument relies on the standard shift-share assumptions: relevance (the instrument predicts actual firm-level delays/shortfalls) and exclusion (the instrument, conditional on controls, only affects the firm's outcomes through its impact on import delays/congestion). That is, port-level congestion patterns from ports a firm historically used should not directly affect the firm's outcomes today, other than via the channel of current import disruptions.

Figure 5 plots the evolution of the average change in firm-level dwell time exposure. The series shows a sharp increase in port congestion during the heightened inflationary period, reflecting severe bottlenecks in global shipping and longer delays in moving goods through US ports.

Figure 6. Marginal Cost Changes: Unit Value and Freight Cost Ratio Exposures
 (A) Mean change in unit cost exposure (B) Mean change in freight cost ratio exposure



Notes: Weighted average of firm-level year-to-year changes in exposure to the unit value and freight cost ratio over time. Panjiva firms. Weights are firm-level annual total imports.

3.3 Marginal Cost Shocks: Unit Import Cost and Freight Cost

We obtain shocks to marginal costs by proxying the imported–input component of marginal cost using two aggregates from U.S. Census merchandise trade: (i) a *unit import cost* index at the HS2 level (indexed by h), $UV_{h,t} \equiv FOB_{h,t}/Q_{h,t}$ (customs value per physical quantity), and (ii) a *freight/insurance* index at the port level (indexed by p), $FR_{p,t} \equiv CIF_{p,t}/FOB_{p,t}$. We take 12-month log changes and map these series to firm f using fixed pre-period (2019) exposure weights derived from Panjiva:

$$\Delta \text{UnitC}_{f,t} = \sum_h \omega_{fh}^{2019} \Delta_{12} \ln UV_{h,t}, \quad \Delta \text{FreightC}_{f,t} = \sum_p \omega_{fp}^{2019} \Delta_{12} \ln FR_{p,t},$$

where ω_{fh}^{2019} is the share of f 's 2019 imports in HS2 category h and ω_{fp}^{2019} is the share from port p .²⁴

Figure 6 shows the evolution of the two data series over time. Both components of import costs—unit values and transportation costs—rose sharply during the inflationary period, reflecting surging global input prices and severe pressures in shipping and logistics.

3.4 Numerator-Panjiva Match

We link the Numerator consumer purchase data to BoL records from Panjiva by matching firm names across the two datasets. The process begins with the manufacturer, parent brand, and brand information in Numerator, and then matches them hierarchically to consignee (and, if needed, shipper names) in Panjiva, after extensive name cleaning and standardization to address name spellings, formats, and inconsistencies (e.g., embedded location details in Panjiva name field). This approach yields a firm-level link between product sales transactions and import activity, which forms the basis for our analysis.

²⁴For exposition, one can also summarize the imported–input contribution as $\Delta MC_{f,t} \approx \phi^c(\Delta \text{UnitC}_{f,t} + \Delta \text{FreightC}_{f,t})$; in practice we use the two components separately.

Details of the matching procedure, data cleaning, and examples are provided in Appendix B.3.

From the matching procedure between our cleaned baseline Numerator product sample and Panjiva data, we obtain the following. Of the 3.5 million product-month observations in Numerator from 2019 to 2023, 2.6 million are from firms whose names match at least once to Panjiva BoL. Among these, 1.5 million observations also have our shortfall measure—that is, they match to the cleaned baseline Panjiva sample of stable importer firms (cleaning procedure described above). We refer to this subset as the baseline Numerator–Panjiva sample, which accounts for 43% of all product-month observations and 50% of total sales. This is our primary sample for the analysis, though complementary samples are used where noted below.

Summary Statistics. The matched Numerator–Panjiva dataset provides a detailed view of how international trade dynamics, such as delivery delays and import disruptions, affect firm-level pricing and sales in the U.S. market. Table 1 reports summary statistics for both the full Numerator sample and the baseline Numerator–Panjiva matched sample.

Table 1. Summary Statistics at the Product-Date Level


	Full Numerator Sample	Numerator-Panjiva Sample
Number of observations	2,906,983	1,251,377
Date range	2019 - 2023	2019 - 2023
Distinct number of products	126,738	48,316
Distinct number of firms	28,767	4,056
Distinct number of brands	43,277	12,333
Distinct number of product-categories	3,982	3,594
Distinct number of product-descriptions	217	211
Average price changes	0.044	0.047
Average sales	1,358,296	1,701,967
Average number of transactions - all users	721.031	956.376
Average number of transactions - static users	342.480	442.228
Panjiva firm dummy	0.430	1.000

Notes: This table presents summary statistics at the product-month level from 2019 to 2023. Column (1) shows statistics for the full Numerator sample (after cleaning procedures as described in Section 3.1). Column (2) shows statistics for the baseline Numerator-Panjiva matched sample (with Panjiva cleaning procedures as described in Section 3.2).


Which firms experienced the largest price increases during the recent inflationary period? Figure 7 shows that bigger firms and those with higher market shares faced the steepest hikes. Panel (A) bin-scatters product-level 12-month price changes against log firm size, defined as average yearly sales in the Numerator, controlling for product department and year effects. Panel (B) plots the same outcome

Figure 7. Price Changes, Firm Size, and Supply Chain Diversification


(A) Product Price Growth vs Firm Size




(B) Product Price Growth vs Firm Market Share



(C) Product Price Growth vs Supplier Countries



(D) Product Price Growth vs Industry-Supplier Countries



Notes: (A) Binscatters of product-level price changes against log firm size, defined as the average yearly sales of the firm in Numerator, residualized for year and department fixed effects. (B) Binscatters of product-level price changes against firm market share, defined as the firm's average sales share in its product category in a year, residualized for year and department fixed effects. (C) Binscatters of product-level price changes against the number of countries the firm sources from, controlling for firm size, residualized for year and department fixed effects. (D) Binscatters of product-level price changes against the number of countries the firm's industry sources from for various HS codes, controlling for firm size, residualized for year and department fixed effects.

against firm market share within product categories, again conditional on department and year effects. The lower panels highlight the role of supply chain diversification. Panel (C) shows a negative correlation between price growth and the number of sourcing countries at the firm level, conditional on firm size. Panel (D) shows a similar negative correlation when diversification is measured at the HS-code level. These raw correlations suggest that firms with broader sourcing networks may have faced fewer disruptions and, in turn, smaller price hikes.

4 Analysis

This section examines how supply-chain disruptions –delivery shortfalls, unit import cost shocks, and freight cost shocks– affect firm pricing. We first estimate pass-through of a firm’s own disruptions, characterize the timing of price adjustments, and document heterogeneity across product categories. We then introduce strategic interactions and show that competitors’ disruptions also raise a firm’s prices—over and above the effect of its own shocks, and that these spillovers extend even to non-importing firms.

4.1 Identifying Pass-Through of Own Supply Chain Disruptions to Prices

Empirical specification. Guided by the simple model’s local decomposition $dp = \alpha dmc + \beta_s d \ln s$, we estimate a reduced form in which (i) availability enters via firm-level delivery shortfalls, our proxy for $\Delta \ln s$, and (ii) marginal cost change is decomposed into unit-import and freight components. Concretely, we estimate

$$\Delta \text{Price}_{p,f,t} = \phi_s \text{Shortfall}_{f,t(k)} + \phi_M \Delta \text{UnitC}_{f,t} + \phi_F \Delta \text{FreightC}_{f,t} + \theta_f + \theta_{j(p)q(t)} + \epsilon_{p,f,t}, \quad (1)$$

where $\Delta \text{Price}_{p,f,t}$ is the 12-month change in the price of product p at firm f in month t ; $\text{Shortfall}_{f,t(k)}$ is the k -month cumulative delivery shortfall (baseline $k = 3$, Section 3.1); and $\Delta \text{UnitC}_{f,t}$, $\Delta \text{FreightC}_{f,t}$ are firm shift-share exposures to unit-import and freight cost changes (Section 3.3). Firm fixed effects (θ_f) absorb time-invariant heterogeneity in pricing, and category $\times$ quarter fixed effects ($\theta_{j(p)q(t)}$) absorb common market-specific demand/supply movements.

Equation (1) is the one-to-one empirical counterpart to the theory, specifically (E3): delivery shortfalls are a monotone proxy for availability, $\Delta \ln s_{f,t} \approx -\phi_1 \text{Shortfall}_{f,t(k)}$ with $\phi_1 > 0$, so the estimated ϕ_s recovers the availability channel scaled by measurement, i.e., $\phi_s = -\beta_s \phi_1$ and is therefore expected to be positive. On the cost side, the model implies that changes in unit-import and freight costs enter dmc with output elasticities θ_M and θ_F ; accordingly, ϕ_M and ϕ_F identify $\alpha \theta_M$ and $\alpha \theta_F$, i.e., the same own pass-through α multiplying the relevant cost elasticities. For comparability with the cost pass-through literature, we scale right-hand-side cost exposures by industry import-intensity weights; shortfalls, similarly, are scaled with intermediates import shares.²⁵

²⁵Specifically, using the BEA Input-Output Accounts data, we compute import shares of 0.094 for those industries that map to consumer product sectors in the Numerator and scale unit cost and freight cost series to align with the interpretation of these coefficients as the pass-through of marginal cost shocks to prices. For shortfall-related series, we scale by 0.125, which is the import share of intermediates, to align with the interpretation of these coefficients as the pass-through of availability shocks to prices.

Idiosyncratic demand shocks may bias OLS estimates of the shortfall coefficient ϕ_s . For example, a positive product-specific demand shock can prompt a firm to accelerate orders, mechanically shrinking measured shortfalls and attenuating ϕ_s toward zero. Measurement error in the shortfall proxy (our local linear map from $\Delta \ln s$ to $\text{Shortfall}_{f,t(k)}$) can generate the same attenuation. To address endogeneity, we implement a shift-share IV that instruments $\text{Shortfall}_{f,t(k)}$ with *Delivery Shortfall Exposure* and *$\Delta \text{Dwell-time Exposure}$* , as described in Section 3.2. These exposures aggregate plausibly exogenous congestion shocks across the firm’s pre-period import portfolio (origins, ports, lanes, and products), providing firm-time variation in predicted shortfalls that is orthogonal to idiosyncratic demand after conditioning on firm and (Numerator) product category $\times$ time effects. This IV design isolates the supply-side component of availability movements, so the second-stage $\hat{\phi}_s$ recovers the model’s availability channel (formally, $-\beta_s \phi_1$), rather than confounding demand.

Results. We begin by estimating baseline pass-through elasticities using both OLS and instrumental variables (IV) to identify the causal effect of supply chain disruptions—delivery shortfalls, import costs, and freight cost shocks—on prices. These estimates provide a benchmark for the magnitude and timing of firms’ price adjustments to these disruptions.

Table 2 reports the results corresponding to equation (1). The dependent variable is the 12-month price change at the product-month level, estimated on the Numerator–Panjiva matched sample for 2020–2023. Our baseline shortfall measure captures deviations from cumulative import deliveries over a three-month horizon ($k = 3$). Column (1) presents OLS estimates, while Column (2) replaces direct shortfalls with a shift-share shortfall exposure measure. Columns (3)–(5) instrument firm-level shortfalls using shortfall exposure and exposure to changes in dwell times. All specifications include firm fixed effects and category-time fixed effects to absorb constant firm-specific pricing factors and demand shifts at the product-category level. Table A.3 shows that the instruments are strong: the weak-IV F-statistics range from 118 to 371, well above conventional thresholds.

Across specifications, all pass-through coefficients on delivery shortfalls and cost shocks are positive, but OLS estimates on shortfalls are biased downward—likely reflecting unobserved demand shocks. Using shift-share exposure in Column (2) raises the estimated effect, while IV estimation further increases the coefficients, yielding stable elasticities above 0.21. The pass-through elasticities on cost shocks are sizable but generally lower than the estimates in the literature.

Columns (4) and (5) extend the specification to include one-month lags of shortfalls and cost variables. For shortfalls, only contemporaneous shocks matter; lagged effects are insignificant, consistent with the shortfall measure already reflecting cumulative three-month deviations. By contrast, for import and freight costs, lagged shocks are more important, suggesting that rising marginal costs pass through

Table 2. Price Effects of Own Supply Chain Disruptions. Baseline Pass-Through Estimates


	ΔP (OLS)	ΔP (OLS-Shift Share)	ΔP (IV)	ΔP (IV)	ΔP (IV)
Shortfall	0.006 (0.0047)	0.108*** (0.0378)	0.211*** (0.0770)	0.241*** (0.0795)	0.267*** (0.0853)
$\Delta UnitC$	0.151** (0.0607)	0.152** (0.0607)	0.146** (0.0613)	0.037 (0.0670)	0.032 (0.0671)
$\Delta FreightC$	0.097** (0.0437)	0.104** (0.0438)	0.154*** (0.0496)	0.028 (0.0585)	0.033 (0.0591)
Lag Shortfall				-0.050 (0.0761)	-0.041 (0.0745)
Lag $\Delta UnitC$				0.214*** (0.0674)	0.218*** (0.0673)
Lag $\Delta FreightC$				0.230*** (0.0552)	0.231*** (0.0552)
Firm FE	✓	✓	✓	✓	✓
Cat-Quarter FE	✓	✓	✓	✓	✓
Observations	969539	969539	968175	939819	939819
Weak IV F-stat			371.986	124.526	118.615

Notes: The table reports regressions of 12-month price changes on measures of own supply chain disruptions, estimated using product-month-level Numerator–Panjiva matched data for 2020–2023. Column (1) uses the *Shortfall* measure in OLS; Column (2) uses the *Shortfall Exposure* measure instead; the remaining columns report IV estimates using *Shortfall Exposure* and *Dwell Time Exposure* as instruments. Column (5) additionally includes an import dummy equal to one if the firm has positive imports in that month. The baseline measure of shortfalls is based on the deviation from cumulative import deliveries over a 3-month horizon ($k = 3$). All specifications include firm and product category–quarter fixed effects. Standard errors, clustered at the product category–quarter level, are reported in parentheses. ***, **, *: significance at the 1%, 5%, 10% levels, respectively.

with a short delay. Column (5) also adds an import dummy (equal to one if the firm imports in that month), but this inclusion does not materially change the results.

We further investigate the timing of firms' price responses to delivery shortfalls in Figure 8. The figure plots coefficients from separate IV regressions of 12-month product-level price changes on delivery shortfalls defined at different horizons of moving average deviations (k months), following the specification in Column 5 of Table 2. The results show that the horizon of the delivery shortfall matters for pass-through. When disruptions are short-lived, firms tend to absorb more of the shocks. As shortfalls become more prolonged, however, pass-through to prices increases, suggesting that firms adjust once they recognize the shocks are persistent and cannot be absorbed internally.

Figure 8. Pass-Through Estimates for Different Horizons of Delivery Shortfalls



Notes: The figure depicts coefficients from separate IV regressions of 12-month product-level price changes on firm-level delivery shortfalls in moving averages at different horizons (in months). The specification follows Column 5 of Table 2 and includes controls for current and lagged changes in unit and freight costs, an import dummy, and firm and product-category-quarter fixed effects. Red dots indicate regression coefficients; dashed lines represent 90% confidence intervals. Standard errors are clustered at the product-category-quarter level.

Overall, the findings show that both disruptions to delivery schedules and increases in marginal costs due to higher import and freight expenses were significantly passed through to consumer prices during the pandemic. Appendix Table A.4 examines heterogeneity by product type. While durable goods experienced markedly higher price growth during the pandemic, evidence of a higher elasticity of delivery-delay pass-through for durables is weak.²⁶ This suggests that much of the observed price growth in durable goods reflected heightened demand rather than stronger supply-side pass-through.

²⁶Following Argente et al. (2020), we measure durability as the inverse of the frequency of trips the average household makes to purchase products in a department over a year. Table A.2 lists the most and least durable departments.

4.2 Strategic Interactions

Extraordinary supply chain disruptions during and after the pandemic, together with an extreme and broad surge in inflation, led many commentators to wonder whether firms passed through aggregate shocks above and beyond the disruptions they directly experienced.²⁷ Yet firms' pricing decisions reflect not only the pass-through of their shocks but also the market forces that govern consumer reallocation and discipline market power. Indeed, a sizable literature documents strategic complementarities in pricing across different settings (Amiti, Itskhoki and Konings, 2019; Albagli et al., 2025), naturally raising the question of whether consumer price increases during this period reflected such strategic interactions among firms. In this section, we extend our baseline analysis of own supply chain disruptions to explicitly incorporate competitive spillovers through rivals' disruptions.

Theoretical motivation. As before, we start by motivating our empirical specifications with theory. Building on the simple model in Section 2, we now allow for strategic interactions across firms. Intuitively, each firm sets prices taking into account not only its own costs and availability but also how rivals' conditions shift the overall competitive environment. Formally, a firm's optimal price can be written as a fixed point in which markups depend on its own fundamentals and on the prevailing sectoral price and availability conditions.²⁸

The key step is to translate this best-response logic into an expression that depends only on fundamentals. Rivals' prices, which appear in the firm's first-order condition, can themselves be expressed in terms of their marginal costs and availability shocks. Aggregating those responses with share weights yields sector-level indices for competitors' marginal costs and availability.²⁹ Substituting these rival responses back into firm i 's condition yields our empirical specification:

$$\Delta p_{it} = \underbrace{\frac{1}{1+\Gamma_{it}}}_{\alpha} \Delta mc_{it} + \underbrace{\frac{\Lambda_{it}}{1+\Gamma_{it}}}_{\beta} \Delta \tau_{it} + \underbrace{\frac{\Gamma_{-it}}{1+\Gamma_{it}}}_{\gamma^{mc}} \bar{\alpha}_{-it} \Delta mc_{-it} + \underbrace{\left(\frac{\Lambda_{-it}}{1+\Gamma_{it}} + \frac{\Gamma_{-it}}{1+\Gamma_{it}} \bar{\beta}_{-it} \right)}_{\delta^{\tau}} \Delta \tau_{-it} + \varepsilon_{it}. \quad (2)$$

The first two coefficients, α and β , correspond to the direct effects from Section 2: pass-through of own costs and the impact of own availability. The new terms capture *strategic interactions*. When competitors' costs rise, their prices adjust upward on average; because firm i 's markup condition depends

²⁷Baqae and Farhi (2022); Glover, Mustre-del Río and von Ende-Becker (2023). The idea of “greedflation” also gained traction, suggesting that firms exploited extraordinary shocks to raise prices to opportunistically increase profits – see Robert Reich’s testimony to Congress (April 5, 2022); Weber and Wasner (2023).

²⁸Formally, firm i 's price satisfies $\tilde{p}_{it} = mc_{it} + \mathcal{M}_i(\tilde{p}_{it}, \tilde{\tau}_{it}, \mathbf{p}_{-it}, \boldsymbol{\tau}_{-it}; \boldsymbol{\xi}_t)$, with Γ_{it} and Λ_{it} denoting the own-price and own-availability derivatives, and Γ_{-it} and Λ_{-it} the sensitivities to rivals. Linearizing and substituting the firm's own price reaction delivers the reduced-form above. Details in Appendix A.2.

²⁹We show in Appendix A that rivals' aggregate responses can be written as $\Delta p_{-it} \approx \bar{\alpha}_{-it} \Delta mc_{-it} + \bar{\beta}_{-it} \Delta \tau_{-it}$, where $\bar{\alpha}_{-it}$ and $\bar{\beta}_{-it}$ are weighted averages of rivals' pass-through and availability sensitivities.

on the sectoral price level, this generates a complementary response, captured by γ^{mc} . Rivals' availability, in turn, influences firm i both directly (more sectoral availability relaxes scarcity, pushing i 's price down) and indirectly through rivals' own price adjustments. The composite coefficient δ^τ therefore summarizes both a direct markup effect and an indirect price-mediated channel. The resulting specification, hence, nests the simple model's scarcity mechanism and quantifies how both own and rival supply-chain fundamentals transmit into prices through the same generalized-elasticity channel.

Empirical specification. Building on these theoretical foundations, we extend the own-effects regression by adding competitor indices of delivery shortfalls and cost movements via unit import and freight cost shocks. Specifically, we estimate:

$$\begin{aligned} \Delta \text{Price}_{p,f,t} = & \phi_s \text{Shortfall}_{f,t(k)} + \phi_M \Delta \text{UnitC}_{f,t} + \phi_F \Delta \text{FreightC}_{f,t} \\ & + \psi_s \text{Shortfall}_{-f,j(p),t(k)} + \psi_M \Delta \text{UnitC}_{-f,j(p),t} + \psi_F \Delta \text{FreightC}_{-f,j(p),t} \\ & + \theta_f + \theta_{j(p)q(t)} + \epsilon_{p,f,t}. \end{aligned} \quad (3)$$

As before, $\Delta \text{Price}_{p,f,t}$ is the 12-month price change for product p at firm f and month t , and the own variables are as in the baseline. Rival indices $X_{-f,j,t}$ are leave-one-out, revenue-share-weighted averages of the rivals' variables— $X \in \{\text{Shortfall}, \Delta \text{UnitC}, \Delta \text{FreightC}\}$ —within market $j = j(p)$. Specifically, with firm revenue shares S_{gjt} for each firm g in market $j(p)$, define $\omega_{fg,t} = S_{gjt}/(1 - S_{fjt})$ and set $X_{-f,j,t} = \sum_{g \neq f} \omega_{fg,t} X_{g,t}$. Firm fixed effects θ_f absorb time-invariant firm heterogeneity in pricing, and category-quarter effects $\theta_{j(p)q(t)}$ absorb common demand or supply shocks. As before, cost exposures and shortfall-related variables are scaled so that coefficients can be read as pass-through elasticities.

4.2.1 Results

Table 3 brings strategic interactions to the fore. As in the baseline, Columns (1)–(2) report OLS and OLS with the shift-share shortfall exposure, while the remaining columns implement the IV strategy that instruments own firm-level shortfalls with shortfall- and dwell-time exposures. Appendix Table A.5 reports the first-stage results, showing strong instruments with F-statistics comfortably above conventional thresholds.

The first three columns tell a familiar story for *own* disruptions— the IV coefficients are sizable and precisely estimated. The novel result is that *competitors'* disruptions also matter for a firm's pricing. In Column (3), the IV estimates imply that delivery delays faced by rivals pass through into a firm's prices with an elasticity around 0.12—smaller than the effect of own shocks, but economically meaningful. Firms also significantly respond to competitors' cost shocks.

Table 3. Pass-Through of Supply Chain Disruptions. Strategic Interactions

	<i>Importing firms</i>			<i>All firms</i>	
	OLS	OLS-Shift Share	IV	IV	IV
Shortfall	0.005 (0.0048)	0.112*** (0.0379)	0.221*** (0.0768)	0.322*** (0.0591)	0.274*** (0.0595)
$\Delta UnitC$	0.155** (0.0608)	0.157*** (0.0608)	0.151** (0.0613)	0.270*** (0.0549)	0.203*** (0.0588)
$\Delta FreightC$	0.098** (0.0438)	0.105** (0.0439)	0.159*** (0.0498)	0.066** (0.0331)	0.088** (0.0414)
Shortfall, compet	-0.003 (0.0132)	-0.007 (0.0130)	0.122*** (0.0464)	0.133*** (0.0286)	
$\Delta UnitC$, compet	0.351** (0.1626)	0.349** (0.1625)	0.346** (0.1637)	0.366*** (0.1314)	
$\Delta FreightC$, compet	0.194* (0.1037)	0.212** (0.1039)	0.234** (0.1053)	0.069 (0.0806)	
Shortfall, compet Imp.					0.108*** (0.0280)
Shortfall, compet Non-Imp					0.117*** (0.0323)
$\Delta UnitC$, compet Imp.					0.616*** (0.1477)
$\Delta UnitC$, compet Non-Imp.					-0.012 (0.1631)
$\Delta FreightC$, compet Imp.					0.019 (0.0875)
$\Delta FreightC$, compet Non-Imp.					0.135 (0.0894)
Firm FE	✓	✓	✓	✓	✓
Cat-Quarter FE	✓	✓	✓	✓	✓
Observations	962815	962815	961451	1671773	1671773
Weak IV F-stat			387.074	705.376	707.173

Notes: The table reports regressions of 12-month price changes on measures of own and competitor's supply chain disruptions, estimated using product-month-level data. Columns (1)-(3) use Numerator–Panjiva matched data for 2020–2023, while columns (4)-(5) also add non-importing firms that do not ever match to Panjiva. Column (1) uses the *Shortfall* measure in OLS; Column (2) uses the *Shortfall Exposure* measure instead; the remaining columns report IV estimates using *Shortfall Exposure* and *Dwell Time Exposure* as instruments for *Shortfall*. All specifications include import dummy and firm- and product category–quarter fixed effects. Standard errors, clustered at the product category–quarter level, are reported in parentheses. ***, **, *. significance at the 1%, 5%, 10% levels, respectively.

Columns (4)–(5) expand the sample to include Numerator firms that never match to Panjiva, pooling importers and non-importers. Column (4) shows that the average pass-through, both own and competitors’, is broadly similar to the importer-only sample. Column (5) then splits the strategic interaction terms by the firm’s importing status. Even non-importing firms—those not directly exposed to supply disruptions—raise prices when their importing competitors are hit, consistent with demand reallocation toward unaffected sellers and the resulting slackening of competitive pressure.³⁰

These findings highlight the role of strategic interactions in shaping the broader impact of supply chain disruptions. Firms not only raised prices following their own supply shocks but also in response to competitors’ disruptions, suggesting that aggregate shocks can propagate through competitive spillovers and amplify overall effects.

Next, we examine how firms’ pass-through of their own and competitors’ delivery delays varies with aggregate conditions. The question is whether pass-through intensifies in certain states of the economy. Table 4 addresses this by re-estimating the baseline specification from column (5) of Table 3, interacting shortfall and competitor-shortfall coefficients with indicators for different states. We consider two dimensions: high-inflation sectors and periods of elevated inflation.

The first column compares high- and low-inflation sectors. A *High-inflation category* is defined as a product category–month with price growth above the annual median across all categories. Pass-through of both own and competitors’ shortfalls is significantly stronger in these sectors. The second column captures periods following high inflation, when inflation perceptions and expectations are likely elevated. A *High-inflation month* is defined as one with previous-month inflation above the annual median. Again, pass-through is higher in these periods.

Overall, the state-dependent results show that inflationary pressures amplify firms’ ability to pass costs through to prices.

5 Quantification

We now turn to a model-consistent accounting exercise in an impulse–response spirit. Because we do not observe clean, disaggregate shocks for the full period, we feed aggregate monthly shock paths into the estimated micro structure and trace how they propagate through market shares and strategic interaction. Concretely, we use three aggregate series: delivery shortfalls, changes in port dwell times, and changes in imported–input costs (unit import prices and freight). Market structure remains fully micro (firm–level revenue shares within product markets), so the exercise maps common shocks into

³⁰We cannot split the *own* shortfall and cost shocks effect by a firm’s importing status because these variables are mechanically zero for non-importers.

Table 4. State-Dependent Pass-Through of Delivery Shortfalls to Prices

	(1) High-inflation category	(2) High-inflation month
Shortfall $\times$ Low State	-0.717*** (0.0731)	0.278*** (0.0579)
Shortfall $\times$ High State	1.167*** (0.0704)	0.367*** (0.0628)
$\Delta UnitC$	0.165*** (0.0633)	0.250*** (0.0555)
$\Delta FreightC$	0.147*** (0.0386)	0.071** (0.0331)
Shortfall, compet $\times$ Low State	-0.175*** (0.0345)	0.112*** (0.0285)
Shortfall, compet $\times$ High State	0.375*** (0.0350)	0.153*** (0.0300)
$\Delta UnitC$, compet	0.460*** (0.1462)	0.339** (0.1319)
$\Delta FreightC$, compet	0.164* (0.0915)	0.073 (0.0805)
Firm FE	✓	✓
Cat-Quarter FE	✓	✓
Observations	1671773	1671773
Weak IV F-stat	351.708	350.956

Notes: The table re-estimates the baseline specification from column (5) of Table 3, interacting shortfall and competitor-shortfall coefficients with indicators for different states. In column (1), a *High-inflation category* is defined as a category-month with price growth above the annual median across all categories. In column (2), a *High-inflation month* is defined as one with previous-month inflation above the annual median. Standard errors, clustered at the product category-quarter level, are reported in parentheses. ***, **, *: significance at the 1%, 5%, 10% levels, respectively.

heterogeneous firm–product price responses via own pass-through and strategic feedback. We obtain a full decomposition into direct and amplified responses to both cost shocks and supply chain disruptions.




Figure 9. Aggregate predicted price change: costs vs. availability, own vs. strategic components

Notes: The black line plots the fitted aggregate price change from the baseline specification. Shaded areas decompose the prediction into channels: *Cost Direct* (own marginal-cost pass-through) and *Cost Amp* (amplification via strategic interactions with rivals’ costs) in blue; *Delay Own* (own delivery shortfalls), *Delay Strat* (rivals’ availability), and their price-mediated amplification components *Amp Own* and *Amp Strat* in green. The decomposition multiplies estimated coefficients from the baseline IV regression by observed firm-level fundamentals and share-weighted competitor indices (leave-one-out weights $S_{jt}/(1-S_{it})$), and then aggregates to the sector level. Series are normalized to the pre-pandemic baseline (2019Q4 = 0); the red dashed line marks the pandemic onset. Positive values denote year-over-year log price changes.

At the core of the exercise is a simple accounting for how monthly product–level prices move. We write the change in log prices as

$$\Delta p_t = \underbrace{(I - \gamma W_t)^{-1}}_{X_t \text{ (strategic multiplier)}} \left[\alpha \Delta mc_t + \beta A_t + \delta W_t A_t \right].$$

Here W_t is the leave–one–out revenue–share matrix that captures who competes with whom, with off-diagonals $w_{ij,t} = S_{jt}/(1-S_{it})$ and zeros on the diagonal; $\gamma > 0$ measures how strongly a firm’s pricing responds to rivals; and $X_t = (I - \gamma W_t)^{-1}$ stacks the rounds of best responses. The bracketed vector collects the *direct* forces: (i) pass-through of marginal-cost changes ($\alpha \Delta mc_t$), (ii) the effect of the firm’s own availability (βA_t), and (iii) scarcity that arrives through rivals, weighted by market structure ($\delta W_t A_t$). The multiplier X_t then turns these direct impulses into equilibrium outcomes by propagating them through the network of competitors. In reporting results, we keep this interpretation front and center: the “direct” piece is the bracket, while “amplification” is the extra movement generated when X_t feeds the shock through rivals and back again.

We use three *aggregate* monthly shock series constructed in Section 3: (i) changes in unit-import

price exposures, (ii) changes in freight cost exposures (our marginal-cost shocks; see Section 3.3), and (iii) an availability composite built from delivery shortfall exposure and port dwell times (see Section 3.2 and Appendix Table A.5 for weights). All series are aligned by month-year, normalized to a 2019Q4 baseline, and (for the baseline specification) expressed in 12-month changes. We then merge these *common* shocks to every product-month in the Numerator panel—there is no product-specific exposure at this stage—so cross-sectional heterogeneity in predicted price responses comes solely from market structure via the leave-one-out revenue-share matrix W_t and the estimated strategic parameters.

Our parameterization reflects the baseline micro estimates in column 4 of Table 3. As a result, we construct aggregate marginal cost by combining unit import and freight cost with their respective output elasticities mirroring (E3). Specifically, we set the output elasticities on unit imports and freight at $\theta_M = 0.27$ and $\theta_F = 0.06$; the own-availability slope at $\beta = 0.322$; the rival-availability slope at $\delta = 0.266$; and the strength of strategic complementarities at $\gamma = 0.54$. The interaction matrix, W_t , uses pre-pandemic revenue shares within narrowly defined product markets.

The resulting decomposition in Figure 9 is intuitive. The fitted aggregate price series is the sum of a cost piece and an availability piece, each with a direct component and an amplification component. Cost shocks—unit values and freight—lift prices on impact via pass-through, and then a bit more as competitors adjust. Availability plays the starring role in the run-up: when deliveries fall or dwell times spike, firms serve fewer arrivals and raise prices; because many firms are hit at once, those increases propagate through the network and are amplified. Dwell times are especially important at the height of congestion, both because they depress availability broadly and because the multiplier is strongest when many rivals face the same bottleneck. The framework also explains spillovers to firms with little direct exposure: even non-importers move prices when they sell into markets where import-reliant rivals set the reference point; so scarcity transmits through competitive interaction.

Two caveats guide interpretation. First, by treating the shortfall and dwell series as common monthly impulses, we assume proportional exposure across firms once we fix market shares; in reality, shocks were highly heterogeneous, but isolating the aggregate and heterogeneous shocks from the data without a fully developed quantitative model is infeasible. Second, we hold parameters and (baseline) market shares fixed, so we do not capture second-round changes in pass-through or market structure. We therefore view the exercise as impulse-response accounting: it gives an indication of the qualitative decomposition of different channels at play, delineating own pass-through of delay and cost shocks, and highlighting the important role strategic complementarity played. Overall, our exercise highlights the importance of delay shocks in combination with strategic complementarity.

6 Conclusion

We study theoretically and empirically how supply chain disruptions shape consumer prices, focusing on both firms' own import shocks and strategic responses to competitors' disruptions. On the theory side, we derive a tractable markup rule that augments the Lerner condition with an endogenous availability term and extend it to a multi-firm setting with strategic complementarities. Empirically, we build a large dataset linking shipment-level Bills of Lading to granular consumer prices and use shift-share instruments based on congestion and input-cost shocks to identify causal pass-through. We find that delivery delays raise prices materially and persistently, while imported-input costs have smaller effects, and that sectoral disruptions spill over across firms, as competitors' shocks also raise prices. An accounting framework shows that these strategic interactions substantially amplify the price impact of supply-chain shocks—most notably delay shocks—during the pandemic.

These findings argue for treating supply-chain resilience as a macro-prudential concern. High-frequency indicators of inbound flows and port congestion could serve as early warnings when availability constraints bind, while targeted investments in port capacity, warehousing, and diversified logistics routings can mitigate the price impact of delays. Because prices are strategic complements, congestion at one set of firms can elevate rivals' markups, underscoring the importance of contestability and competition in dampening amplification.

A natural next step is embedding our tractable pricing-inventory block in a full GE framework with endogenous sourcing and inventories. Such a model would capture how firms' collective responses—shifting suppliers or transport modes and holding precautionary inventories—can themselves deepen congestion and amplify shocks in equilibrium. This would allow to quantify the aggregate effects of dis-aggregate supply chain disruptions and the evaluation of policy counterfactuals, such as port expansions, routing diversification, or greater competition at logistics nodes.

References

- Aguirregabiria, Victor.** 1999. “The Dynamics of Markups and Inventories in Retailing Firms.” *The Review of Economic Studies*, 66(2): 275–308.
- Albagli, Elías, Francesco Grigoli, Emiliano Luttini, Dagoberto Quevedo, and Marco Rojas.** 2025. “Beyond Costs: The Dominant Role of Strategic Complementarities in Pricing.”
- Alessandria, George, Joseph P. Kaboski, and Virgiliu Midrigan.** 2010. “Inventories, Lumpy Trade, and Large Devaluations.” *American Economic Review*, 100(5): 2304–2339.
- Alessandria, George, Shafaat Yar Khan, Armen Khederlarian, Carter Mix, and Kim J. Ruhl.** 2023. “The Aggregate Effects of Global and Local Supply Chain Disruptions: 2020–2022.” *Journal of International Economics*, 146: 103788.
- Alvarez, Luis J., Emmanuel Dhyne, Marco Hoeberichts, Claudia Kwapil, Hervé Le Bihan, Patrick Lünnemann, Fernando Martins, Roberto Sabbatini, Harald Stahl, Philip Vermeulen, and Jouko Vilmunen.** 2006. “Sticky Prices in the Euro Area: A Summary of New Micro-Evidence.” *Journal of the European Economic Association*, 4(2–3): 575–584.
- Amiti, Mary, Oleg Itskhoki, and Jozef Konings.** 2019. “International Shocks, Variable Markups, and Domestic Prices.” *The Review of Economic Studies*, 86(6): 2356–2402.
- Amiti, Mary, Sebastian Heise, Fatih Karahan, and Ayşegül Şahin.** 2024. “Inflation Strikes Back: The Role of Import Competition and the Labor Market.” *NBER Macroeconomics Annual*, 38: 71–131.
- Argente, David, Salomé Baslandze, Douglas Hanley, and Sara Moreira.** 2020. “Patents to Products: Product Innovation and Firm Dynamics.” Federal Reserve Bank of Atlanta FRB Atlanta Working Paper 2020-4.
- Ascari, Guido, Dennis Bonam, and Andra Smadu.** 2024. “Global supply chain pressures, inflation, and implications for monetary policy.” *Journal of International Money and Finance*, 142: 103029.
- Bai, Xiwen, Jesús Fernández-Villaverde, Yiliang Li, and Francesco Zanetti.** 2024. “The causal effects of global supply chain disruptions on macroeconomic outcomes: Evidence and theory.” *SSRN Electron. J.*
- Baqae, David, and Emmanuel Farhi.** 2022. “Supply and demand in disaggregated Keynesian economies with an application to the COVID-19 crisis.” *Am. Econ. Rev.*, 112(5): 1397–1436.
- Baslandze, Salomé, Simon Fuchs, KC Pringle, and Michael Sparks.** 2025. “Tariffs and Consumer Prices: Insights from Newly Matched Consumption-Trade Micro Data.” *Policy Hub*, 2025(1): 1–17.
- Borusyak, Kirill, and Xavier Jaravel.** 2021. “The Distributional Effects of Trade: Theory and Evidence from the United States.” <https://www.nber.org/papers/w28957>.
- Brancaccio, Giulia, Myrto Kalouptsi, and Theodore Papageorgiou.** 2024. “Investment in Infrastructure and Trade: The Case of Ports.” , (w32503).
- Brancaccio, Giulia, Myrto Kalouptsi, and Theodore Papageorgiou.** 2025. “Rigidities in transportation and supply chain disruptions.” *AEA Pap. Proc.*, 115: 543–550.

- Carreras-Valle, Maria-Jose, and Alessandro Ferrari.** 2025. "The Cost of Delivery Delays." *AEA Papers and Proceedings*, 115: 618–623.
- Cavallo, Alberto, and Oleksiy Kryvtsov.** 2023. "What can stockouts tell us about inflation? Evidence from online micro data." *J. Int. Econ.*, 103769.
- Comin, Diego A, Robert C Johnson, and Callum J Jones.** 2023. "Supply Chain Constraints and Inflation." <https://www.nber.org/papers>, , (31179).
- Davis, Steven, John Haltiwanger, and Scott Schuh.** 1998. *Job Creation and Destruction*. Vol. 1. 1 ed., The MIT Press.
- De Loecker, Jan, Pinelopi K Goldberg, Amit K Khandelwal, and Nina Pavcnik.** 2016. "Prices, markups, and trade reform." *Econometrica*, 84(2): 445–510.
- Finck, David, Mathias Klein, and Peter Tillmann.** 2024. "The Inflationary Effects of Global Supply Chain Shocks: Evidence from Swedish Microdata."
- Flaaen, Aaron, Flora Haberkorn, Logan Lewis, Anderson Monken, Justin Pierce, Rosemary Rhodes, and Madeleine Yi.** 2023. "Bill of lading data in international trade research with an application to the COVID-19 pandemic." *Review of International Economics*, 31(3): 1146–1172.
- Fuchs, Simon, and Woan Foong Wong.** 2022. "Multimodal Transport Networks." 2022(13).
- Glover, Andrew, José Mustre-del Río, and Alice von Ende-Becker.** 2023. "How much have record corporate profits contributed to recent inflation?" *Econ. Rev.*, 108(1): 1–13.
- Hacıoğlu Hoke, Sinem, Leo Feler, and Jack Chylak.** 2024. "A Better Way of Understanding the US Consumer: Decomposing Retail Sales by Household Income."
- He, Zijian, and Yang Su.** 2023. "Measuring Welfare Gains from Online Stores:Theory and Evidence from the Supreme Court's Wayfair Decision." *SSRN Electronic Journal*.
- Hristakeva, Sylvia, Jura Liaukonyte, and Leo Feler.** 2024. "The no-hunger games: How GLP-1 medication adoption is changing consumer food purchases." *Social Science Research Network*.
- Iglehart, Donald L, and Ward Whitt.** 1970. "Multiple channel queues in heavy traffic. II: sequences, networks, and batches." *Adv. Appl. Probab.*, 2(2): 355–369.
- Khan, Aubhik, and Julia K. Thomas.** 2007. "Inventories and the Business Cycle: An Equilibrium Analysis of (S, s) Policies." *American Economic Review*, 97(4): 1165–1188.
- Kingman, J F C.** 1961. "The single server queue in heavy traffic." *Math. Proc. Camb. Philos. Soc.*, 57(4): 902–904.
- Kryvtsov, Oleksiy, and Virgiliu Midrigan.** 2013. "Inventories, Markups, and Real Rigidities in Menu Cost Models." *Review of Economic Studies*, 80(1): 249–276.
- Lafrogne-Joussier, Raphael, Julien Martin, and Isabelle Mejean.** 2023. "Energy cost pass-through and the rise of inflation: Evidence from French manufacturing firms." C.E.P.R. Discussion Papers CEPR Discussion Papers 18596.

- Liu, Ernest, Vladimir Smirnyagin, and Aleh Tsyvinski.** 2024. "Supply chain disruptions and supplier capital in US firms." *Available at SSRN 4866869*.
- Meyer, Brent H., Brian C. Prescott, and Xuguang Simon Sheng.** 2023. "The impact of supply chain disruptions on business expectations during the pandemic." *Energy Economics*, 126: 106951.
- Nakamura, Emi, and Jón Steinsson.** 2008. "Five Facts About Prices: A Reevaluation of Menu Cost Models." *The Quarterly Journal of Economics*, 123(4): 1415–1464.
- Weber, Isabella M, and Evan Wasner.** 2023. "Sellers' inflation, profits and conflict: why can large firms hike prices in an emergency?" *Rev. Keynes. Econ.*, 11(2): 183–213.

A Additional Derivations

This appendix collects formal derivations for Section 2.1 and Section 4.2, as well as additional discussions and microfoundations.

A.1 Derivations: Simple Model (Section 2.1)

We retain the primitives and notation used in the text. Prices and marginal costs are parameterized in *logs*, p and c , with corresponding *levels* $P := e^p$ and $MC := e^c$. Demand intensity is $\lambda(p) = \Lambda e^{-\sigma p}$ with $\Lambda, \sigma > 0$. Replenishment lead times are i.i.d. exponential with rate $\mu > 0$ (mean $1/\mu$). The firm follows a base-stock policy at level $\tau \in \mathbb{N}$ (lost sales). Define

$$r \equiv \frac{\lambda}{\lambda + \mu} \in (0, 1), \quad s(\tau, \lambda, \mu) = 1 - r^\tau, \quad \kappa(\tau, \lambda, \mu) \equiv \lambda \frac{\partial s}{\partial \lambda} = -\tau(1 - r)r^\tau.$$

Accepted (served) demand arrives at rate λs . One convenient closed form for expected on-hand inventory is

$$\mathbb{E}[I(\tau, \lambda, \mu)] = \frac{\tau - r(\tau + 1) + r^{\tau+1}}{1 - r}. \quad (\text{A1})$$

Lemma 1 (Profit-rate equivalence). *Let demand be Poisson with intensity $\lambda(p) = \Lambda e^{-\sigma p}$ and lead times i.i.d. $\text{Exp}(\mu)$, independent of demand. Under a base-stock policy at level τ with lost sales, stationary/ergodic inventory, instantaneous sales at demand epochs when $I(t) > 0$, instantaneous delivery resets to τ , and a constant per-unit level margin $(P - MC)$, the long-run average contribution (revenue net of unit cost, excluding holding cost) equals margin times the throughput of accepted demand:*

$$\lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \int_0^T (P - MC) \mathbf{1}\{I(t) > 0\} dN(t) = (P - MC) \lambda s(\tau, \lambda, \mu) = (P - MC) \lambda (1 - r^\tau). \quad (\text{A23})$$

Equivalently, in renewal cycles delimited by consecutive deliveries,

$$\frac{(P - MC) \mathbb{E}[\min\{N, \tau\}]}{\mathbb{E}[W]} = (P - MC) \lambda (1 - r^\tau), \quad (\text{A24})$$

where $W \sim \text{Exp}(\mu)$ is the cycle length, $N \mid W \sim \text{Pois}(\lambda W)$, and hence $N \sim \text{Geom}(\frac{\mu}{\lambda + \mu})$ on $\{0, 1, 2, \dots\}$. The full average profit rate including holding cost is

$$\Pi(p, \tau) = (P - MC) \lambda (1 - r^\tau) - h \mathbb{E}[I(\tau, \lambda, \mu)], \quad \mathbb{E}[I] \text{ given by (A1)}. \quad (\text{A25})$$

Proof of Lemma 1 (renewal-reward). Consider renewal cycles starting immediately after a delivery (on-hand = τ) and ending at the next delivery. Let $W \sim \text{Exp}(\mu)$ be the cycle length. Conditional on W , the number of demand arrivals is $N \mid W \sim \text{Pois}(\lambda W)$. Integrating out W (Poisson-exponential mixture) yields a geometric law on $\{0, 1, 2, \dots\}$ with success probability $p := \mu/(\lambda + \mu)$, so $\mathbb{P}\{N = n\} = (1 - p)^n p$ and $1 - p = r = \lambda/(\lambda + \mu)$. Sales in a cycle equal $\min\{N, \tau\}$. Using the truncated-geometric identity,

$$\mathbb{E}[\min\{N, \tau\}] = \sum_{n=0}^{\tau-1} n(1 - p)^n p + \tau(1 - p)^\tau = \frac{1 - p}{p} [1 - (1 - p)^\tau] = \frac{\lambda}{\mu} (1 - r^\tau).$$

By the renewal-reward theorem,

$$\text{avg. sales rate} = \frac{\mathbb{E}[\min\{N, \tau\}]}{\mathbb{E}[W]} = \mu \cdot \frac{\lambda}{\mu} (1 - r^\tau) = \lambda (1 - r^\tau).$$

Multiplying by $(P - MC)$ yields (A23). Subtracting the holding-cost flow $h \mathbb{E}[I]$ gives (A25). $\square^1$

Remark. The acceptance fraction used here is the *arrival-epoch* probability $s(\tau, \lambda, \mu) = 1 - r^\tau$; by PASTA it equals the time-average availability, which is the appropriate object for throughput and profit-rate calculations.

Proof of Proposition 1. Profits per unit time in *levels* are

$$\Pi(p, \tau) = (P - MC) \lambda(p) s(\tau, \lambda(p), \mu) - h \mathbb{E}[I(\tau, \lambda(p), \mu)], \quad P := e^p, \quad MC := e^c. \quad (\text{A2})$$

Step 1 (pricing FOC for given τ). Let $Q(p) := \lambda(p) s(\tau, \lambda(p), \mu)$. Then

$$\frac{\partial}{\partial p} [(P - MC) Q(p)] = P Q(p) + (P - MC) Q'(p) = 0.$$

Divide by $P Q(p) > 0$:

$$\frac{P - MC}{P} = -\frac{Q'(p)}{Q(p)} = -\frac{d \ln Q}{dp} = -\frac{d \ln \lambda}{dp} - \frac{\partial \ln s}{\partial \ln \lambda} \frac{d \ln \lambda}{dp}.$$

Since $d \ln \lambda / dp = -\sigma$ and $\partial \ln s / \partial \ln \lambda = (\lambda/s) \partial s / \partial \lambda = \kappa/s$, we obtain the generalized Lerner rule,

$$\frac{P^* - MC}{P^*} = \frac{1}{\sigma(1 + \kappa/s)}, \quad (\text{A3})$$

equivalently

$$p^* = c - \ln \left(1 - \frac{1}{\sigma(1 + \kappa/s)} \right). \quad (\text{A4})$$

Existence/uniqueness. The LHS of (A3), $1 - e^{c-p}$, is strictly increasing in p from 0 (at $p = c$) toward 1; the RHS is weakly decreasing in p because higher p lowers λ , raises s , and pushes $\kappa/s \uparrow 0$, increasing the denominator. Hence a unique solution exists.

Step 2 (discrete τ FOC). Using

$$\Delta s(\tau) = s(\tau+1) - s(\tau) = r^\tau(1-r), \quad (\text{A5})$$

and, from (A1),

$$\Delta \mathbb{E}[I] = \mathbb{E}[I]_{|\tau+1} - \mathbb{E}[I]_{|\tau} = 1 - r^{\tau+1}, \quad (\text{A6})$$

the exact discrete condition equates marginal revenue and marginal holding cost:

$$(P^* - MC) \lambda r^{\tau^*} (1-r) = h(1 - r^{\tau^*+1}). \quad (\text{A7})$$

Step 3 (continuous relaxation of τ). Let $x := r^{\tau^*} \in (0, 1)$. Using $\partial_\tau s = -x \ln r$ and $\partial_\tau \mathbb{E}[I] = 1 + \frac{r x \ln r}{1-r}$, the FOC becomes

$$(P^* - MC) \lambda (-\ln r) x = h \left(1 + \frac{r x \ln r}{1-r} \right). \quad (\text{A8})$$

This linear equation in x has the unique solution

$$r^{\tau^*} = x^* = \frac{h}{(-\ln r) \left((P^* - MC) \lambda + \frac{h r}{1-r} \right)} \in (0, 1). \quad (\text{A9})$$

¹*Alternative proof (PASTA).* Let $N(t)$ be the demand counting process (rate λ) and $L(t) := \mathbf{1}\{I(t) > 0\}$ the availability indicator. Instantaneous contribution is $(P - MC) L(t) dN(t)$. By PASTA and stationarity,

$$\lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \int_0^T (P - MC) L(t) dN(t) = (P - MC) \lambda \mathbb{E}[L(0)] = (P - MC) \lambda \mathbb{P}\{I > 0\}.$$

Considering the superposed event process of demands and deliveries (i.i.d. types with $\mathbb{P}(\text{demand}) = r$, $\mathbb{P}(\text{delivery}) = 1-r$), an arrival is served iff not all of the immediately preceding τ events were demands, which occurs with probability $1 - r^\tau$. Hence $\mathbb{P}\{I > 0\} = 1 - r^\tau$ and we recover $(P - MC) \lambda (1 - r^\tau)$. Including holding cost reproduces (A25). $\square$

□

Proof of Corollary 1. We establish the availability elasticity, the sign of the price response via the implicit function theorem, the inventory response, and the limits.

Step 1 (availability elasticity). With $s = 1 - r^\tau$ and $\partial_{\ln \mu} r = -r(1 - r)$,

$$\partial_{\ln \mu} s = \frac{\partial s}{\partial r} \partial_{\ln \mu} r = (-\tau r^{\tau-1})(-r(1 - r)) = \tau r^\tau (1 - r),$$

so

$$\partial_{\ln \mu} \ln s = \frac{\tau r^\tau (1 - r)}{1 - r^\tau} > 0. \quad (\text{A11})$$

Step 2 (price response). Let $z := \kappa/s \in (-1, 0)$ and recall the levels Lerner rule (A3). Define

$$\phi(p, \mu) := \underbrace{1 - e^{c-p}}_{=(P-MC)/P} - \frac{1}{\sigma(1+z(p, \mu))} = 0. \quad (\text{A12})$$

Holding μ fixed, $\partial z / \partial p > 0$; holding p fixed, $\partial z / \partial \ln \mu > 0$. Hence

$$\phi_p = e^{c-p} + \frac{1}{\sigma} \frac{1}{(1+z)^2} \frac{\partial z}{\partial p} > 0, \quad \phi_{\ln \mu} = \frac{1}{\sigma} \frac{1}{(1+z)^2} \frac{\partial z}{\partial \ln \mu} > 0. \quad (\text{A13})$$

By the implicit function theorem,

$$\partial_{\ln \mu} P^* = -\frac{\phi_{\ln \mu}}{\phi_p} < 0. \quad (\text{A14})$$

Step 3 (inventory response). Use the discrete FOC (A7):

$$F(\tau, \mu) := (P^* - MC) \lambda r^\tau (1 - r) - h(1 - r^{\tau+1}) = 0. \quad (\text{A15})$$

Differentiate $F(\tau^*(\mu), \mu) = 0$ w.r.t. $\ln \mu$: $F_\tau \partial_{\ln \mu} \tau^* + F_{\ln \mu} = 0$, i.e.,

$$\partial_{\ln \mu} \tau^* = -\frac{F_{\ln \mu}}{F_\tau}. \quad (\text{A16})$$

We have

$$F_\tau = \ln r [(P^* - MC) \lambda r^\tau (1 - r) + h r^{\tau+1}] < 0$$

since $\ln r < 0$ and the bracket is positive; and

$$\begin{aligned} F_{\ln \mu} &= (P^* - MC) \lambda \partial_{\ln \mu} [r^\tau (1 - r)] + \lambda r^\tau (1 - r) \partial_{\ln \mu} P^* - h \partial_{\ln \mu} (1 - r^{\tau+1}) \\ &= h(1 - r^{\tau+1}) [r - \tau(1 - r)] + \lambda r^\tau (1 - r) \partial_{\ln \mu} P^* - h(\tau + 1) r^{\tau+1} (1 - r) \leq 0, \end{aligned}$$

where we used $(P^* - MC) \lambda = \frac{h(1 - r^{\tau+1})}{r^\tau (1 - r)}$, $\partial_{\ln \mu} [r^\tau (1 - r)] = r^\tau (1 - r) [r - \tau(1 - r)]$, $\partial_{\ln \mu} (1 - r^{\tau+1}) = (\tau + 1) r^{\tau+1} (1 - r)$, and $\partial_{\ln \mu} P^* = P^* \partial_{\ln \mu} P^* < 0$ from (A14). Thus $F_{\ln \mu} < 0$ and, with $F_\tau < 0$, we conclude

$$\partial_{\ln \mu} \tau^* < 0. \quad (\text{A17})$$

Step 4 (limits). As $\mu \rightarrow \infty$, $r \rightarrow 0$, hence $s \rightarrow 1$ and $z \rightarrow 0^-$; (A3) implies

$$\frac{P^* - MC}{P^*} \rightarrow \frac{1}{\sigma} \Rightarrow P^* \rightarrow MC \cdot \frac{\sigma}{\sigma - 1} \quad \text{and} \quad p^* \rightarrow c + \ln \frac{\sigma}{\sigma - 1}.$$

As $\mu \rightarrow 0$, $1 - r \sim \mu/(\lambda + \mu)$, so $s = 1 - r^\tau \sim \tau(1 - r) \sim \tau \mu/(\lambda + \mu)$; moreover $\partial_{\ln \mu} z \rightarrow 0$ so the pricing response becomes locally flat, $\partial_{\ln \mu} P^* \rightarrow 0^-$. □

From generalized Lerner to a linear price rule. Start from the generalized markup condition

$$p - c = \mathcal{M}(p, \ln s) := \frac{1}{\sigma(1 + z(p, \ln s))}, \quad z := \kappa/s \in (-1, 0). \quad (\text{D1})$$

Define $F(p, c, \ln s) := p - c - \mathcal{M}(p, \ln s) = 0$. Totally differentiating around an operating point and collecting terms gives

$$(1 - \mathcal{M}_p) dp = dc + \mathcal{M}_{\ln s} d \ln s, \quad (\text{D2})$$

so the linear price rule is

$$dp = \alpha dc + \beta_s d \ln s + \varepsilon, \quad \alpha := \frac{1}{1 + \Gamma} \in (0, 1), \quad \beta_s := \frac{\Lambda}{1 + \Gamma} < 0, \quad (\text{D3})$$

with $\Gamma := -\mathcal{M}_p > 0$ and $\Lambda := \mathcal{M}_{\ln s} < 0$, all evaluated at the operating point. For $\mathcal{M}(p, \ln s) = [\sigma(1 + z)]^{-1}$,

$$\mathcal{M}_p = -\frac{z_p}{\sigma(1 + z)^2}, \quad \mathcal{M}_{\ln s} = -\frac{z_{\ln s}}{\sigma(1 + z)^2} \Rightarrow \Gamma = \frac{z_p}{\sigma(1 + z)^2} > 0, \quad \Lambda = -\frac{z_{\ln s}}{\sigma(1 + z)^2} < 0. \quad (\text{D4})$$

Interpretation: a higher price relaxes scarcity ($z_p > 0$) so own pass-through $\alpha \in (0, 1)$; higher availability raises the effective elasticity ($z_{\ln s} > 0$) so the availability coefficient $\beta_s < 0$. The specification $dp = \alpha dc + \beta_s d \ln s$ is the estimable reduced form used in the empirics.

A.2 Derivations: Strategic Interactions (Section 4.2)

A.2.1 Deriving the Markup and Stock-in Best Response

Consider the profit maximization problem of firm i written in the conjectural variation form, where the firm chooses both its log price p_i and its log stock-in shifter τ_i . The problem is given by

$$\begin{aligned} \max_{p_i, \tau_i, \mathbf{p}_{-i}, \boldsymbol{\tau}_{-i}} \quad & \left\{ \exp\{p_i + q_i(p_i, \mathbf{p}_{-i}, \tau_i, \boldsymbol{\tau}_{-i}; \boldsymbol{\xi}) + s_i(q_i, \tau_i)\} \right. \\ & \left. - TC_i(\exp\{q_i(p_i, \mathbf{p}_{-i}, \tau_i, \boldsymbol{\tau}_{-i}; \boldsymbol{\xi})\}, \tau_i) \right\} \\ \text{subject to} \quad & h_{-i}(p_i, \tau_i, \mathbf{p}_{-i}, \boldsymbol{\tau}_{-i}; \boldsymbol{\xi}) = 0. \end{aligned}$$

Here, p_i and q_i are the log price and log quantity demanded of firm i ; $s_i(q_i, \tau_i)$ is a function that shifts the effective stock-in probability (so that higher s_i implies improved product availability); $TC_i(\cdot, \tau_i)$ is the total cost function (in levels); and $h_{-i}(\cdot)$ is the conjectural variation vector function with elements $h_{ij}(\cdot)$ for $j \neq i$. (For brevity we omit the time subscript t .)

This formulation nests models of monopolistic competition, oligopolistic Bertrand competition, and oligopolistic Cournot competition, provided that the demand system is invertible. In particular, to capture firm behavior under monopolistic and oligopolistic Bertrand competition we may choose the conjectural variation function as

$$h_{-i}(p_i, \tau_i, \mathbf{p}_{-i}, \boldsymbol{\tau}_{-i}; \boldsymbol{\xi}) = \begin{cases} \mathbf{p}_{-i} - \mathbf{p}_{-i}^*, \\ \boldsymbol{\tau}_{-i} - \boldsymbol{\tau}_{-i}^*, \end{cases}$$

which corresponds to the assumption that firm i believes its choices do not affect its competitors' prices (set at $\mathbf{p}_{-i}^*$) or stock-in shifters (set at $\boldsymbol{\tau}_{-i}^*$). The case of Cournot competition requires choosing $h_{-i}(\cdot)$ so that it implies $\mathbf{q}_{-i} \equiv \mathbf{q}_{-i}^*$ for some given $\mathbf{q}_{-i}^*$; with an invertible demand system, this can be ensured by setting

$$h_{-i}(p_i, \tau_i, \mathbf{p}_{-i}, \boldsymbol{\tau}_{-i}; \boldsymbol{\xi}) = -(\mathbf{q}_{-i}(p_i, \mathbf{p}_{-i}, \tau_i, \boldsymbol{\tau}_{-i}; \boldsymbol{\xi}) - \mathbf{q}_{-i}^*).$$

Thus, the firm's behavior under competition in both prices and product availability is captured by a conditional profit maximization problem.

We introduce the following notation:

1. $e^{p_i + q_i} \lambda_{ij}$ for $j \neq i$ denotes the Lagrange multipliers associated with the constraints in the maximization problem.
2. $\zeta_{ijk}(\mathbf{p}, \boldsymbol{\xi}) \equiv \partial h_{ij}(\mathbf{p}, \boldsymbol{\xi}) / \partial p_k$ is the elasticity of the conjectural variation function with respect to p_k , with $\zeta_{ijj}(\cdot) > 0$ as a normalization; the matrix $Z_i \equiv \{\zeta_{ijk}(\cdot)\}_{j,k \neq i}$ is assumed to have full rank. Similarly, we define

$$\theta_{ijk}(\mathbf{p}, \boldsymbol{\xi}) \equiv \frac{\partial h_{ij}(\mathbf{p}, \boldsymbol{\xi})}{\partial \tau_k},$$

and denote by $\Theta_i \equiv \{\theta_{ijk}(\cdot)\}_{j,k \neq i}$ the corresponding matrix of cross-stock-in elasticities.

3. The own-price elasticity of demand is

$$\epsilon_i(\mathbf{p}, \boldsymbol{\xi}) \equiv -\frac{\partial q_i(\mathbf{p}, \boldsymbol{\xi})}{\partial p_i} > 0,$$

and the cross-price elasticities are given by

$$\delta_{ij}(\mathbf{p}, \boldsymbol{\xi}) \equiv \frac{\partial q_i(\mathbf{p}, \boldsymbol{\xi})}{\partial p_j} \quad (j \neq i).$$

4. Similarly, define the effective own-stock-in elasticity as

$$\eta_i(\mathbf{p}, \boldsymbol{\tau}, \boldsymbol{\xi}) \equiv \frac{\partial (q_i + s_i)}{\partial \tau_i} > 0,$$

and the cross-stock-in elasticities as

$$\gamma_{ij}(\mathbf{p}, \tau, \xi) \equiv \frac{\partial(q_i + s_i)}{\partial \tau_j} \quad (j \neq i).$$

The first-order conditions (after simplification) for the profit maximization problem with respect to p_i are then given by

$$\begin{aligned} (1 - \epsilon_i + \epsilon_i e^{-\mu_i}) + \sum_{k \neq i} \lambda_{ik} \zeta_{iki} &= 0, \\ \forall j \neq i \quad (-\delta_{ij} + \delta_{ij} e^{-\mu_i}) + \sum_{k \neq i} \lambda_{ik} \zeta_{ikj} &= 0, \end{aligned}$$

where the log markup is defined by $\mu_i \equiv p_i - mc_i$, with $mc_i \equiv \ln(\partial TC_i / \partial Q_i)$ being the log marginal cost. In the presence of the stock-in shifter, the effective demand depends on both q_i and s_i ; consequently, the perceived elasticity of demand is modified and given by

$$\sigma_i \equiv \epsilon_i(1 + \kappa_i) - \zeta_i' Z_i^{-1} \delta_i - \theta_i' \Theta_i^{-1} \gamma_i,$$

where $\kappa_i \equiv \partial s_i / \partial q_i$, and $\zeta_i \equiv \{\zeta_{iji}\}_{j \neq i}$ and $\delta_i \equiv \{\delta_{ij}\}_{j \neq i}$ are $(N-1) \times 1$ vectors. The matrix $Z_i \equiv \{\zeta_{ijk}\}_{j \neq i, k \neq i}$ has full rank by the invertibility assumption, and similarly for the stock-in part with $\theta_i \equiv \{\theta_{iji}\}_{j \neq i}$ and $\gamma_i \equiv \{\gamma_{ij}\}_{j \neq i}$ and the matrix $\Theta_i \equiv \{\theta_{ijk}\}_{j \neq i, k \neq i}$.

Solving the first-order conditions yields the expression for the optimal markup of firm i :

$$\mu_i = \ln \frac{\sigma_i}{\sigma_i - 1}.$$

Since ϵ_i , ζ_{ijk} , and δ_{ij} are functions of $(\mathbf{p}, \xi)$ (and now also of τ through s_i), it follows that $\sigma_i \equiv \sigma_i(\mathbf{p}, \tau; \xi)$. Therefore, we define the log markup function as

$$\mathcal{M}_i(\mathbf{p}, \tau; \xi) \equiv \ln \frac{\sigma_i(\mathbf{p}, \tau; \xi)}{\sigma_i(\mathbf{p}, \tau; \xi) - 1}.$$

Thus, the optimal price of firm i solves the fixed point equation

$$\tilde{p}_i = mc_i + \mathcal{M}_i(\tilde{p}_i, \tilde{\tau}_i, \mathbf{p}_{-i}, \tau_{-i}; \xi).$$

An analogous argument, based on the first-order condition with respect to the stock-in shifter τ_i , leads to a similar fixed point characterization for the optimal stock-in decision. Defining the effective stock-in elasticity as

$$\phi_i \equiv \eta_i - \zeta_i^{\tau'} Z_i^{-1} \delta_i^{\tau} - \theta_i^{\tau'} \Theta_i^{-1} \gamma_i,$$

where $\eta_i \equiv \partial(q_i + s_i) / \partial \tau_i$ is the own elasticity with respect to τ_i and $\delta_i^{\tau} \equiv \{\partial q_i / \partial \tau_j\}_{j \neq i}$, and where ζ_i^{τ} and θ_i^{τ} are the corresponding derivatives of the conjectural variation function with respect to τ_i , we define the stock-in premium function by

$$\mathcal{T}_i(p_i, \tau_i, \mathbf{p}_{-i}, \tau_{-i}; \xi) \equiv -\ln\left(1 - \frac{mc_i^{\tau}}{\phi_i}\right),$$

with mc_i^{τ} the log marginal cost of adjusting the stock-in shifter. Hence, the optimal stock-in shifter satisfies the fixed point equation

$$\tilde{\tau}_i = mc_i^{\tau} + \mathcal{T}_i(\tilde{p}_i, \tilde{\tau}_i, \mathbf{p}_{-i}, \tau_{-i}; \xi).$$

This completes the proof of Proposition 1 (Extended). Notice that when the stock-in shifter is absent (or when $s_i(\cdot)$ is constant), the above derivation reduces to the standard framework in which only price is chosen. In our extended model, the generalized elasticities σ_i and ϕ_i capture both the direct demand sensitivities and the strategic interdependencies across firms in the dimensions of price and product availability.

A.2.2 Deriving the Exposure Indices

Assume that product availability enters the expenditure function via a stock-in adjustment factor. In particular, let the log aggregate expenditure be given by

$$z_t = \ln \min_{\{Q_{it}\}} \left\{ \sum_{i=1}^N P_{it} Q_{it} s_{it}(\tau_{it}) \mid U(\{Q_{it}\}; Q_t) = 1 \right\}.$$

Here, $s_{it}(\tau_{it})$ adjusts the effective expenditure for product i based on its stock-in probability (or availability), with τ_{it} being the firm's stock-in shifter.

Let $E(s)$ denote the minimized expenditure:

$$E(s) = \min_{\{Q_{it}\}} \left\{ \sum_{i=1}^N P_{it} Q_{it} s_{it}(\tau_{it}) \mid U(\{Q_{it}\}; Q_t) = 1 \right\}$$

so that $z_t = \ln E(s)$. At the optimum, let Q_{it}^* be the optimal quantity for firm i . Define the Lagrangian for the cost minimization problem as

$$L(\{Q_{it}\}, \lambda; s) = \sum_{i=1}^N P_{it} Q_{it} s_{it}(\tau_{it}) - \lambda (U(\{Q_{it}\}; Q_t) - 1)$$

By the Envelope Theorem, the derivative of the minimized expenditure $E(s)$ with respect to s_{it} is given by the direct partial derivative evaluated at the optimum, i.e.,

$$\frac{\partial E(s)}{\partial s_{it}} = \frac{\partial L}{\partial s_{it}} \Big|_{Q_{it}=Q_{it}^*} = P_{it} Q_{it}^*$$

(Any indirect effects via the optimal Q_{it}^* vanish because the first-order conditions for Q_{it} ensure that $\partial L / \partial Q_{it} = 0$.) Taking the derivative of $z_t = \ln E(s)$ with respect to s_{it} , we obtain

$$\frac{\partial z_t}{\partial s_{it}} = \frac{1}{E(s)} \cdot \frac{\partial E(s)}{\partial s_{it}} = \frac{P_{it} Q_{it}^*}{\sum_{k=1}^N P_{kt} Q_{kt}^*}$$

In summary, by the Envelope Theorem applied to this cost minimization problem, differentiating z_t with respect to s_{it} gives

$$\frac{\partial z_t}{\partial s_{it}} = \frac{P_{it} Q_{it}^*}{\sum_{k=1}^N P_{kt} Q_{kt}^*} \equiv S_{it},$$

where Q_{it}^* is the optimal quantity for firm i and S_{it} is its expenditure (or revenue) market share. Now, if the availability adjustment is given by a function $s_{it} = s(\tau_{it})$, then by the chain rule we have

$$\frac{\partial z_t}{\partial \tau_{it}} = \frac{\partial z_t}{\partial s_{it}} \cdot \frac{\partial s(\tau_{it})}{\partial \tau_{it}} = S_{it} \frac{\partial s(\tau_{it})}{\partial \tau_{it}}.$$

Next, suppose the markup function for firm i depends on competitors' stock-in shifters only through an aggregate availability index, so that we can write

$$\mathcal{M}_i(p_t, \tau_t; \xi_t) = \tilde{\mathcal{M}}_i(p_t, z(s(\tau_t)); \xi_t).$$

Then, by the chain rule the partial derivative of the markup function with respect to a competitor j 's stock-in shifter is

$$\frac{\partial \mathcal{M}_i}{\partial \tau_j} = \frac{\partial \tilde{\mathcal{M}}_i}{\partial z} \cdot \frac{\partial z}{\partial s(\tau)} \cdot \frac{\partial s(\tau)}{\partial \tau_j}.$$

By the Envelope condition we have already shown that

$$\frac{\partial z}{\partial s(\tau)} \cdot \frac{\partial s(\tau)}{\partial \tau_j} = S_{jt}.$$

Hence, it follows that

$$\frac{\partial \mathcal{M}_i}{\partial \tau_j} = \frac{\partial \tilde{\mathcal{M}}_i}{\partial z} S_{jt}.$$

When constructing the competitor delay index, we define the weight on competitor j as

$$\omega_{ijt}^\tau \equiv \frac{\frac{\partial \mathcal{M}_i}{\partial \tau_j}}{\sum_{k \neq i} \frac{\partial \mathcal{M}_i}{\partial \tau_k}} = \frac{\frac{\partial \tilde{\mathcal{M}}_i}{\partial z} S_{jt}}{\frac{\partial \tilde{\mathcal{M}}_i}{\partial z} \sum_{k \neq i} S_{kt}} = \frac{S_{jt}}{1 - S_{it}},$$

since the market shares of competitors sum to $1 - S_{it}$. Consequently, the competitor delay index is given by

$$d\tau_{-it} \equiv \sum_{j \neq i} \omega_{ijt}^\tau d\tau_{jt} = \sum_{j \neq i} \frac{S_{jt}}{1 - S_{it}} d\tau_{jt}.$$

This completes the proof: by explicitly incorporating product availability via the function $s(\tau_{it})$ into the aggregate expenditure function, we derive—using the Envelope condition—that the marginal effect of a competitor's stock-in shifter on the markup function is proportional to its market share. Therefore, the competitor delay index aggregates individual delay changes with weights $\omega_{ijt}^\tau = S_{jt}/(1 - S_{it})$, analogous to the competitor price index in the original framework.

A.3 Additional Derivations

A.3.1 Microfoundation for exponential lead times.

In the simple model we take the delivery (replenishment) lead time W to be exponential with rate μ . A standard way to rationalize both the *shape* (memorylessness) and the *scale* of W is via a heavy-traffic approximation for an upstream GI/GI/1 queue that processes replenishment orders.²

Setup. Consider an FCFS GI/GI/1 system that represents the upstream fulfillment “server.” Interarrival times have mean $1/\lambda$ and squared coefficient of variation c_a^2 ; service times S have mean $\mathbb{E}[S] = 1/\mu_s$ and squared coefficient of variation c_s^2 . Let the traffic intensity be $\rho := \lambda \mathbb{E}[S] \in (0, 1)$. For a tagged order submitted in steady state, let W denote its waiting (lead) time.

Lemma 2 (Heavy-traffic exponentiality of lead times). *As $\rho \uparrow 1$ (heavy traffic), the stationary lead time is asymptotically exponential after the canonical scaling:*

$$\frac{(1-\rho)}{\mathbb{E}[S]} W \stackrel{d}{\rightarrow} \text{Exp}\left(\frac{2}{c_a^2 + c_s^2}\right).$$

Equivalently, for ρ close to one the distribution of W is well approximated by $\text{Exp}(\mu_{\text{HT}})$ with

$$\mu_{\text{HT}} \approx \frac{2(1-\rho)}{(c_a^2 + c_s^2)\mathbb{E}[S]}, \quad \mathbb{E}[W] \approx \frac{(c_a^2 + c_s^2)}{2} \frac{\mathbb{E}[S]}{1-\rho}.$$

Proof sketch. The stationary workload (and hence the waiting time under FCFS) of a GI/GI/1 queue admits a diffusion approximation in heavy traffic: the centered-scaled workload process converges to a reflected Brownian motion (RBM) with negative drift $\theta \propto (1-\rho)/\mathbb{E}[S]$ and variance parameter proportional to $c_a^2 + c_s^2$. The stationary distribution of an RBM is exponential with rate $2\theta/\sigma^2$, yielding the claimed limit for the scaled W . Matching the mean with Kingman’s formula, $\mathbb{E}[W] \approx \frac{\rho}{1-\rho} \cdot \frac{c_a^2 + c_s^2}{2} \mathbb{E}[S]$, confirms the scale. $\square$

Interpretation and mapping to the simple model. Lemma 2 provides a structural justification for taking lead times to be exponential: in congested upstream environments the *entire distribution* of W is approximately memoryless, with a rate that moves one-for-one with the slack $1-\rho$ and is dampened by variability ($c_a^2 + c_s^2$). In our toy model we write $W \sim \text{Exp}(\mu)$ and treat μ as a reduced-form “replenishment speed.” The lemma implies that, under a GI/GI/1 microfoundation,

$$\mu \equiv \mu_{\text{HT}} \approx \frac{2(1-\rho)}{(c_a^2 + c_s^2)\mathbb{E}[S]},$$

so shocks that raise upstream congestion ($\rho \uparrow$) or variability (c_a^2 or $c_s^2 \uparrow$) lower μ and lengthen expected lead times $1/\mu$. From the firm’s perspective—especially when it is small relative to the sector— μ is naturally taken as exogenous, summarizing the state of the upstream queue; the exponential assumption concerns the *shape*, while μ collects congestion and volatility into a single sufficient statistic.

Sanity checks. (i) In the special case of an $M/M/1$ system ($c_a^2 = c_s^2 = 1$), the heavy-traffic mean reduces to $\mathbb{E}[W] \approx \mathbb{E}[S]/(1-\rho)$, consistent with the exact $M/M/1$ formula at $\rho \uparrow 1$. (ii) Lower variability or more slack (smaller ρ) raises μ_{HT} and thus makes the exponential approximation concentrate at shorter lead times, matching the comparative statics we exploit in the toy model’s pricing and inventory results.

A.3.2 Comparison with Kryvtsov and Midrigan (2013)

Pricing rules in compact form. Kryvtsov–Midrigan (KM) derive a frictionless *reset* price that multiplies a static markup by an intertemporal shadow-value term:

$$P_i(s^t) = \underbrace{\frac{\varepsilon_i(s^t)}{\varepsilon_i(s^t) - 1}}_{\text{static markup}} \times \underbrace{(1 - \delta_z) E_t \left[\frac{Q(s^{t+1})}{Q(s^t)} \Omega(s^{t+1}) \right]}_{\text{value of a marginal unit in inventory}},$$

²Classical references include Kingman (1961) and Iglehart and Whitt (1970). We use the first-order stationary heavy-traffic limit and Kingman’s formula as consistency checks for the mean.

where the *elasticity of expected sales* $\varepsilon_i(s^t)$ rises with availability (fewer stockouts), compressing markups, and the intertemporal term prices the shadow value of inventories. In our static base-stock model, availability enters *algebraically* through the levels Lerner rule (Appendix A.1, eq. (A3)):

$$\frac{P^* - MC}{P^*} = \frac{1}{\sigma(1+z)}, \quad z := \kappa/s \in (-1, 0),$$

so the effective demand elasticity is $\sigma(1+z)$. The inventory margin is summarized by the exact discrete stockout-avoidance condition (Appendix A.1, eq. (A7)):

$$(P^* - MC) \lambda r^\tau (1-r) = h(1 - r^{\tau+1}),$$

which equates the marginal revenue from an extra unit of buffer to its holding cost.

Alignment and differences. (i) *Elasticity channel.* KM's $\varepsilon_i(s^t)$ and our $\sigma(1+z)$ parameterize the same mechanism: better availability raises the effective elasticity and lowers markups; both coincide with the frictionless benchmark as stockouts vanish ($z \rightarrow 0^-$).

(ii) *Intertemporal vs. static scarcity.* KM price the shadow value of inventories with a discounted continuation term; in our steady-state, risk-neutral setting this collapses to the one-period stockout-avoidance margin in (A7). Scarcity in both frameworks ultimately operates through the elasticity object that governs the markup.

(iii) *Units and mapping.* KM's level markup $\varepsilon/(\varepsilon - 1)$ maps to our gross markup $\mu = P/MC = \sigma(1+z)/[\sigma(1+z) - 1]$ (from (A3)); locally, their reset-price rule linearizes to the same inverse-elasticity logic that underpins our log markup $m = p - c = -\ln(1 - 1/[\sigma(1+z)])$.

(iv) *Scope.* KM embed inventories in a dynamic, nominally sticky environment where reset prices track discounted targets. We collapse the same forces to closed-form static objects, which yields sharp pass-through coefficients and a linear, empirically tractable decomposition into cost and availability channels used in the main text.

B Data Appendix

B.1 Numerator Data

Numerator collects data from households in several ways. Using a mobile phone app called “Receipt Hog,” consumers can (1) snap and upload a picture of their paper receipts, (2) allow Numerator to scrape their emails for digital receipts, and (3) link loyalty and membership accounts (such as Amazon, UberEats, Walmart, and Home Depot accounts), which Numerator then scrapes for transaction information. Panelists are rewarded with coins redeemable for Amazon or Visa gift cards or for cash through PayPal. On average, Numerator rewards panelists approximately \$10 per month for providing their purchase information and completing surveys ([Hacioğlu Hoke, Feler and Chylak, 2024](#))

Our baseline sample excludes the items whose Numerator sector is either “Limited Service Restaurant,” “Restaurant,” “Non-Item,” “Unknown,” or “Indeterminate Category.” In addition, we restrict the sample to the items for which we have non-missing brand or manufacturer information and at least one non-missing product classification from *category*, *majorcat*, or *department* or a detailed product description. These restrictions give us a sample of 83,352,791 unique items sold during the period 2019-2023.³ For these items, we have information on their classification: sub-category, category, major category, department, and sector (listed hierarchically from more detailed to more aggregated classification). An example here would be the “Adult Cough/Chest Congestion” sub-category in the “Adult Cold Cough & Flu” category, “Cold, Cough & Flu” major category, “Personal Health Care” department, and the “Health & Beauty” sector. For more than 20 million items in our sample, we have detailed product descriptions that provide more information about the packaging, weight, and other product attributes (e.g., CV MUCUS REL 30 CT). In our regression analysis, we use a benchmark classification consisting of 4,000 distinct product categories defined by the *product category* in Numerator, or the broader *major category* description if the category description is missing.

Table [A.1](#) shows the number of stores across various distribution channels in the data, while Figure [A.1](#) shows expenditure shares and the number of unique item IDs by broadly defined sectors in Numerator.

³Item ID is the finest product code available in Numerator. It is the product code that the data provider assigns based on what looks like distinct products, however because of the lack of detailed product descriptions and barcodes for many items, it is possible that exactly the same product is sold in different stores, but since it lacks barcode information, it has to be treated as different item by the Numerator. In our analysis, since we define products at the firm-brand-product category level, all these distinct item IDs that in reality describe the same product will be pooled together, getting closer to the definition of a product.

Table A.1. Stores: Summary Statistic of Distribution Channel

Channel	Frequency	Percent	Cum.	Channel	Frequency	Percent	Cum.
Food	10,142	20.85%	20.85%	Military	49	0.10%	98.01%
FSR - Regional/Ethnic	8,428	17.32%	38.17%	School Tuition & Fees	49	0.10%	98.11%
Gas & Convenience	7,421	15.25%	53.43%	FSR - American	46	0.09%	98.21%
Liquor	2,905	5.97%	59.40%	Sports & Recreation	46	0.09%	98.30%
Bodega	2,317	4.76%	64.16%	Mortgage Payment	44	0.09%	98.39%
Apparel	1,692	3.48%	67.64%	FSR - Midscale	41	0.08%	98.48%
Drug	1,471	3.02%	70.66%	Farm	36	0.07%	98.55%
LSR - Bakery/Cafe	1,270	2.61%	73.27%	Movie Theatre	36	0.07%	98.62%
Beauty	999	2.05%	75.33%	Outlet Store	36	0.07%	98.70%
Pet	990	2.04%	77.36%	Sports Entertainment	36	0.07%	98.77%
Home Improvement	819	1.68%	79.05%	Other Association Fees	32	0.07%	98.84%
Craft	701	1.44%	80.49%	Wireless	32	0.07%	98.90%
Online	531	1.09%	81.58%	Charities	31	0.06%	98.97%
Other	515	1.06%	82.64%	Video	31	0.06%	99.03%
LSR - Ethnic/Regional	512	1.05%	83.69%	FSR - Italian/Pizza	29	0.06%	99.09%
Dispensaries	433	0.89%	84.58%	FSR - Seafood/Steak	28	0.06%	99.15%
Sporting Goods Stores	424	0.87%	85.45%	LSR - Salad/Healthful	28	0.06%	99.21%
Postal Services	401	0.82%	86.27%	Transport Hub	28	0.06%	99.26%
Other Retail Store	390	0.80%	87.08%	Music Stores	27	0.06%	99.32%
LSR - Coffee/Bakery	385	0.79%	87.87%	Gambling	24	0.05%	99.37%
Book	357	0.73%	88.60%	Professional Services	24	0.05%	99.42%
Other Specialty Store	331	0.68%	89.28%	Veterinarians	24	0.05%	99.47%
Dollar	311	0.64%	89.92%	Public Markets	21	0.04%	99.51%
Specialty Food Retailer	300	0.62%	90.54%	Yard Services	19	0.04%	99.55%
Mass	237	0.49%	91.03%	Banks	17	0.03%	99.58%
Health	226	0.46%	91.49%	Doctors Office	17	0.03%	99.62%
Electronics	219	0.45%	91.94%	Wholesale	17	0.03%	99.65%
Shoe	203	0.42%	92.36%	Club	15	0.03%	99.69%
Baby & Toy	201	0.41%	92.77%	Parking Lot or Garage	15	0.03%	99.72%
Auto	194	0.40%	93.17%	Laundromat	13	0.03%	99.74%
LSR - Pizza	192	0.39%	93.56%	Airline	11	0.02%	99.77%
Department Store	181	0.37%	93.94%	Other Travel	11	0.02%	99.79%
FSR - Miscellaneous	181	0.37%	94.31%	Telecom	10	0.02%	99.81%
LSR - Burger	137	0.28%	94.59%	Utility Company	10	0.02%	99.83%
Discount Store	134	0.28%	94.87%	Media	9	0.02%	99.85%
Home Furnishings	128	0.26%	95.13%	Real Estate Services	9	0.02%	99.87%
Vapor Stores	121	0.25%	95.38%	Concert Hall or Theater	8	0.02%	99.88%
FSR - Sports Bar	106	0.22%	95.59%	Other Government Payments	8	0.02%	99.90%
Tobacco Shops	97	0.20%	95.79%	Child Care	7	0.01%	99.91%
Auto Services	95	0.20%	95.99%	Retail Services	7	0.01%	99.93%
Other Entertainment	89	0.18%	96.17%	Church Offerings	5	0.01%	99.94%
Other Services	85	0.17%	96.35%	Car Rental	4	0.01%	99.95%
Office	81	0.17%	96.51%	Copy Centers	4	0.01%	99.95%
Shopping Centers & Malls	74	0.15%	96.67%	Public Storage	3	0.01%	99.96%
Amusement Parks	71	0.15%	96.81%	Security Services	3	0.01%	99.97%
Spas	69	0.14%	96.95%	Taxi or Limousine	3	0.01%	99.97%
LSR - Dessert Snack	67	0.14%	97.09%	Travel Agency	3	0.01%	99.98%
Hotels & Resorts	66	0.14%	97.23%	Bakery	2	0.00%	99.98%
CloseOut	62	0.13%	97.35%	Cruise Line	2	0.00%	99.99%
Healthcare	61	0.13%	97.48%	Dance & Comedy Clubs	2	0.00%	99.99%
Florists	53	0.11%	97.59%	Bar or Night Club	1	0.00%	99.99%
LSR - Chicken	53	0.11%	97.70%	Movers	1	0.00%	100.00%
LSR - Mexican	52	0.11%	97.80%	Ticket Outlet	1	0.00%	100.00%
LSR - Sandwich/Deli	52	0.11%	97.91%	Unknown	1	0.00%	100.00%
Total	47,631	97.91%		Total	48,648	100.00%	


Notes: The table shows the number of retailers identifying with a given channel. The sample is the entire Numerator data (2018-2023) with 48,648 unique retailers and 108 unique channels. [A.1](#)

Table A.2. Top 20 and Bottom 20 Departments by Durability

Department	Sector	Durability Index	Department	Sector	Durability Index
Architecture & Design	Books	1.000	Paper & Plastic	Household	0.017
Birthday	Toys	1.000	Lottery	Non-Item	0.017
Business Office Furniture	Office	1.000	LSR Beverages	LSR	0.017
Office Lighting	Office	1.000	Packaged Bakery (Bread & Alternative)	Grocery	0.016
Posters	Books	1.000	Canned	Grocery	0.016
Sustainability	Books	1.000	Shelf Stable Meals	Grocery	0.016
Celebrate Childrens Books	Books	0.975	Alcohol Beverages	Grocery	0.014
Miscellaneous (Books)	Books	0.975	Baking & Cooking	Grocery	0.012
Parenting & Families	Books	0.966	Combustible Nicotine Products	Tobacco	0.011
Pressure & Temperature	Tools & Home Imp.	0.959	Candy (Snacks)	Grocery	0.011
Performing Arts	Books	0.945	Meat	Grocery	0.009
Romance	Books	0.911	Infant Toddler Nutrition	Baby	0.009
Carrying Cases	Office	0.910	Frozen Foods	Grocery	0.008
Biography & Memoirs	Books	0.909	Snack	Grocery	0.006
Office Furniture & Lighting	Office	0.896	Pet Food & Treats	Pet	0.006
Stationary (Baby)	Baby	0.890	Dairy	Grocery	0.005
Pregnancy & Maternity	Baby	0.859	Beverages	Grocery	0.005
Sports & Recreation (Books)	Books	0.836	Mobile App Downloads	Electronics	0.005
Travel & Nature (Books)	Books	0.833	Produce	Grocery	0.003

Notes: The table reports the 20 most and 20 least durable departments, where *durability* is defined as the inverse of the average number of transactions per user within a department. The calculation is based on users who were continuously present in the static panel for all 12 months of 2023. Tools & Home Imp: Tools and Home Improvement. LSR: Limited Service Restaurant. Tobacco: Tobacco Products and Accessories.

Figure A.1. Consumer Expenditures and Products by Sectors in Numerator




Notes: (A) The figure shows the product-sector level share of total Numerator expenditures (\$4 billion from \$683 million transactions) in 2022. A rough benchmark comparison with the 2022 Consumer Expenditure Survey (CEX) suggests that the relatively large share of grocery-related spending in Numerator is expected. The CEX describes 52.6% of all expenditures as belonging to categories similar to those captured by Numerator. Among these similar categories, food comprised 69.1% of spending, while in Numerator, a similar category (Grocery, FSR (Full-Service Restaurants), LSR (Limited-Service Restaurants), Tobacco) comprised 51.4% of consumer outlays. (B) This figure shows the number of unique UPCs by product-sector from the set of all 2022 transactions from our baseline sample.

Figure A.2. Demographic Comparison between Numerator and ACS



Notes: The US Census ACS data is from IPUMS, weighted by personal weight (perwt), and covers 2019. Demographic information from Numerator data is based on a static sample of 400,000 panelists covering the period 2019-2022. For this comparison, we do not apply the Numerator demographic weights (*demo_weight*) that are designed to balance the sample further to make it representative of the US Census household characteristics.

Figure A.3. Sectoral Heterogeneity in Price Changes



Notes: Distribution of median 12-month product-level price changes within each product category by quarter derived from the Numerator data.

B.2 BoL Panjiva Data

We construct the dataset using BoL records on U.S. imports from 2010 to 2023, taking *conname* as the importer name and aggregating its shipments to measure imports.⁴ We keep only shipments with the U.S. as the final destination⁵, retain only single-HS code shipments⁶, and drop records with redacted or missing (cleaned) firm names, using the cleaning procedure described below.

We then aggregate the data to a firm-HS2-month panel by summing shipment volumes in TEU. To ensure comparability with Numerator data, we restrict the sample to 2018–2023, balance the panel to the minimum and maximum months for each firm-HS2 pair, and drop 2018 after constructing moving cumulative measures. For each firm-HS2-month, we compute 2-, 3-, 4-, 5-, 6-, 9-, and 12-month moving sums of shipment volumes.

To convert shipment volumes to values, we merge BoL data with U.S. Trade Online data at the HS2 level, applying 2019 TEU-to-KG conversion ratios. The monthly Census data from U.S. Trade Online report containerized vessel values and weights (in kg) by 2-digit HS code. We first compute monthly unit prices for each HS2 code in 2019, along with conversion factors between weights (Census) and volumes (BoL) by dividing weight by volume for each HS2. These unit prices and conversion factors are then applied to BoL monthly import volumes at the firm-HS2 level to estimate monthly import values. Appendix Figure A.4 compares aggregate BoL import values with maritime U.S. imports from Census. The two series coincide in 2019 by construction and track each other closely over time, with BoL showing some underestimation in 2022.⁷


⁴Some importers are retailers rather than manufacturers. We verify that retailer imports are not conflated with those of third-party companies selling products through the retail chain. Inspection of the Panjiva *notifyparty* field indicates that shipments imported directly by retailers pertain to their own operations—whether for resale under private labels or for other needs. When third-party sellers ship to a retailer, the BoL typically lists the manufacturer or seller as *conname* and the retailer as *notifyparty*, often with the retailer’s name and address.

⁵Some shipments arrive at a U.S. port (e.g., Seattle) but are then forwarded to another country, such as Canada; we exclude these.

⁶Shipments with multiple HS codes cannot be split reliably across codes and account for only a small share of the sample.


⁷This likely reflects the use of 2019 conversion factors for BoL import values, whereas import costs spiked in 2022.

Figure A.4. US Imports over Time. Panjiva vs. Census




Notes: Aggregate imports in BoL Panjiva against USA Trade Online - U.S. Census Bureau imports data. Panjiva import volumes are indexed to the 2019 unit value of an HS code in Census data.

Figure A.5. US Imports over Time. Panjiva vs. Census. Panjiva Firms Importing in 2019.



Notes: Total imports of a sample of firms with positive deliveries in 2019 against USA Trade Online - U.S. Census Bureau imports data. Panjiva import volumes are indexed to the 2019 unit value of an HS code in the Census data.

Figure A.6. Sectoral Heterogeneity in Delivery Shortfalls



Notes: Delivery shortfalls constructed as mean level deviations in the same month compared to the 2019 baseline, HS2 categories weighted with firm-level value shares.

B.3 Numerator-BoL Panjiva Match

We match the Numerator and BoL Panjiva datasets using company names. In BoL, we primarily use consignee names (*conname*) as the firm identifier, which often represent manufacturing companies but can also include retailers. In some cases, U.S. multinational firms ship products directly from abroad to U.S. locations, so their names may also appear as shippers in BoL. We retain the shipper names for our matching procedure, too.

In Numerator, each product entry can be associated with multiple identifiers—manufacturer, parent brand, and brand names—providing several avenues for linking to producers. Manufacturer names are often missing from scanned receipts, and the data provider prioritizes supplementing this information for high-sales products. As a result, larger firms are more likely to have manufacturer data, while it is often absent for smaller firms or less significant products. Our sample contains 9,569 unique manufacturer names, 70,205 unique parent brand names, and 85,903 unique brand names. Examples include:

- PROCTER & GAMBLE COMPANY / REGENERIST OLAY / OLAY
- PRIVATE LABEL / PRIVATE LABEL / WALMART
- LEGO SYSTEMS, INC / DISNEY LEGO / LEGO
- APPLE INC. / APPLE / APPLE WATCH SERIES 8

These manufacturer–parent brand–brand combinations provide a strong basis for matching with firm names in Panjiva. Numerator’s list of roughly 27,000 store names also helps identify Panjiva shipments received by retail stores.

We clean and standardize names from both datasets to ensure consistency, building on routines such as those implemented in [Argente et al. \(2020\)](#) and adding steps tailored to Panjiva’s idiosyncrasies. This includes removing location details (city, country, zip code) frequently embedded in consignee names (e.g., “LG ELECTRONICS PANAMA S.A.” or “TARGET ATLANTA GEORGIA 30309 USA”).

Matching begins from the Numerator, where each observation includes the manufacturer, parent brand, and brand names. We perform exact matches with Panjiva consignee and shipper names using a hierarchical approach: first matching manufacturer names, then parent brands, and finally brand names to consignee names. If no match is found, we repeat the process using shipper names. This procedure yields a comprehensive link between the two datasets, enabling the analysis of firm-level price changes in relation to import activity. The matched Numerator–Panjiva sample relies on this link. For unmatched firms, we still use manufacturer and parent brand names to define firm identifiers, but drop observations where only the brand name is known to reduce noise.

C Additional Empirical Results

Table A.3. Price Effects of Own Supply Chain Disruptions. I-stage Estimates

	Col(3) I-stage Shortfall	Col(4) I-stage Shortfall Lag Shortfall		Col(5) I-stage Shortfall Lag Shortfall	
Shortfall exposure	0.479*** (0.0178)	0.483*** (0.0162)	-0.066*** (0.0164)	0.441*** (0.0160)	-0.099*** (0.0162)
$\Delta Dwell$	-0.163*** (0.0145)	-0.153*** (0.0145)	-0.078*** (0.0156)	-0.134*** (0.0142)	-0.063*** (0.0156)
$\Delta UnitC$	-0.056** (0.0257)	-0.029 (0.0254)	-0.021 (0.0255)	-0.044* (0.0252)	-0.033 (0.0253)
$\Delta FreightC$	0.258*** (0.0210)	0.253*** (0.0209)	0.038* (0.0204)	0.250*** (0.0206)	0.036* (0.0204)
Lag Shortfall exposure		0.022 (0.0159)	0.521*** (0.0168)	0.042*** (0.0157)	0.536*** (0.0166)
Lag $\Delta Dwell$		-0.008 (0.0125)	-0.096*** (0.0136)	-0.021* (0.0123)	-0.106*** (0.0136)
Lag $\Delta UnitC$		-0.039 (0.0246)	-0.087*** (0.0257)	-0.021 (0.0244)	-0.072*** (0.0256)
Lag $\Delta FreightC$		0.020 (0.0218)	0.139*** (0.0228)	0.028 (0.0216)	0.145*** (0.0227)
Observations		939819	939819	939819	939819
F-test of excl. instruments	371.99	229.73	274.32	193.20	300.44

Notes: The table reports I-stage estimates for columns (3)-(5) in Table 2. Standard errors, clustered at the product category–quarter level, are reported in parentheses. ***, **, *: significance at the 1%, 5%, 10% levels, respectively.

Table A.4. Price Effects of Own Supply Chain Disruptions: By Product Durability

	Δp		
	Benchmark IV	Durability Index	Durability Dummy
Shortfall	0.272*** (0.0822)	0.116 (0.1221)	0.261*** (0.0803)
Shortfall $\times$ Durability Index		3.197 (2.5267)	
Shortfall $\times$ Durability Dummy			0.074 (0.0993)
$\Delta UnitC$	0.032 (0.0671)	0.027 (0.0674)	0.030 (0.0672)
$\Delta FreightC$	0.032 (0.0590)	0.004 (0.0606)	0.032 (0.0587)
Lag $\Delta UnitC$	0.220*** (0.0671)	0.223*** (0.0672)	0.219*** (0.0672)
Lag $\Delta FreightC$	0.226*** (0.0547)	0.223*** (0.0548)	0.227*** (0.0547)
Firm FE	✓	✓	✓
Cat-Quarter FE	✓	✓	✓
Observations	939819	939819	939819
Weak IV F-stat	337.650	20.994	196.871

Notes: The table reports regressions analogous to the benchmark IV own-pass-through estimate in column (5) of Table 2, but with shortfalls interacted with the *Durability Index* (column 2) and the *Durability Dummy* (column 3). Product durability is constructed based on the product purchase frequency, following [Argente et al. \(2020\)](#). Standard errors, clustered at the product category–quarter level, are reported in parentheses. ***, **, *: significance at the 1%, 5%, 10% levels, respectively.

Table A.5. Strategic Interactions. I stage

	Col(3) I-stage Shortfall	Col(4) I-stage Shortfall	Col(5) I-stage Shortfall
Shortfall exposure	0.480*** (0.0176)	0.381*** (0.0156)	0.374*** (0.0156)
$\Delta Dwell$	0.171*** (0.0138)	0.343*** (0.0101)	0.352*** (0.0104)
$\Delta UnitC$	0.052** (0.0245)	0.022 (0.0235)	0.053** (0.0242)
$\Delta FreightC$	-0.263*** (0.0198)	-0.371*** (0.0108)	-0.309*** (0.0159)
Shortfall, compet	-0.584*** (0.0077)	-0.451*** (0.0062)	
$\Delta UnitC$, compet	0.054 (0.0645)	0.068 (0.0476)	
$\Delta FreightC$, compet	-0.117** (0.0455)	-0.084*** (0.0326)	
Shortfall, compet Imp.			-0.433*** (0.0066)
Shortfall, compet Non-Imp.			-0.482*** (0.0067)
$\Delta UnitC$, compet Imp.			-0.041 (0.0563)
$\Delta UnitC$, compet Non-Imp.			0.244*** (0.0489)
$\Delta FreightC$, compet Imp.			-0.156*** (0.0365)
$\Delta FreightC$, compet Non-Imp.			0.001 (0.0339)
Observations	961451	1671773	1671773
F-test of excl. instruments	387.07	705.38	707.17

Notes: The table reports I-stage estimates for columns (3)-(5) in Table 3. Standard errors, clustered at the product category–quarter level, are reported in parentheses. ***, **, *: significance at the 1%, 5%, 10% levels, respectively.